

**This microfiche was
produced according to
ANSI/AIIM Standards
and meets the
quality specifications
contained therein. A
poor blowback image
is the result of the
characteristics of the
original document.**

NASA Conference Publication 3320



Human-Computer Interaction and Virtual Environments

Compiled by
Ahmed K. Noor

ADDITIONAL INFORMATION

100-62

100-

1040 0052

Proceedings of a workshop sponsored by the
National Aeronautics and Space Administration,
Washington, D.C., and the University of Virginia Center
for Computational Structures Technology, Hampton,
Virginia, and held at Virginia Consortium of Engineering
and Science Universities, Hampton, Virginia
April 26-27, 1995

November 1995



NASA Conference Publication 3320

Human-Computer Interaction and Virtual Environments

Compiled by

Ahmed K. Noor

University of Virginia Center for Computational Structures Technology • Hampton, Virginia

Proceedings of a workshop sponsored by the
National Aeronautics and Space Administration,
Washington, D.C., and the University of Virginia Center
for Computational Structures Technology, Hampton,
Virginia, and held at Virginia Consortium of Engineering
and Science Universities, Hampton, Virginia
April 26-27, 1995

National Aeronautics and Space Administration
Langley Research Center • Hampton, Virginia 23681-0001

November 1995

This publication is available from the following sources:

NASA Center for AeroSpace Information
800 Elkridge Landing Road
Linthicum Heights, MD 21090-2934
(301) 621-0390

National Technical Information Service (NTIS)
5285 Port Royal Road
Springfield, VA 22161-2171
(703) 487-4650

PREFACE

This document contains the proceedings of the Workshop on Human-Computer Interaction and Virtual Environments held in Hampton, Virginia, April 26-27, 1995. The workshop was jointly sponsored by the University of Virginia Center for Computational Structures Technology and NASA. Workshop attendees came from government agencies, energy laboratories, industry and universities. The objectives of the workshop were to assess the state-of-technology and level of maturity of several areas in human-computer interaction and to provide guidelines for focused future research leading to effective use of these facilities in the design/fabrication and operation of future high-performance engineering systems.

Certain materials and products are identified in this publication in order to specify adequately the materials and products that were investigated in the research effort. In no case does such identification imply recommendation or endorsement of products by NASA, nor does it imply that the materials and products are the only ones or the best ones available for this purpose. In many cases equivalent materials and products are available and would probably produce equivalent results.

Ahmed K. Noor
Center for Computational Structures Technology
University of Virginia
Hampton, VA

CONTENTS

PREFACE	iii
ATTENDEES	vii
INTRODUCTION	xi
HIGHLIGHTS OF THE WORKSHOP	1
Ahmed K. Noor	
THE DESIGN OF PICTORIAL INSTRUMENTS	13 - 1
Stephen Ellis	
ALTERNATIVE DISPLAY AND INTERACTION DEVICES	25 - 2
M. T. Bolas, I. E. McDowall, R. X. Mead, E. R. Lorimer, J. E. Hackbush and C. Greuel	
COMPUTATIONAL VIRTUAL REALITY (VR) AS A HUMAN-COMPUTER INTERFACE IN THE OPERATION OF TELEROBOTIC SYSTEMS	37 - 3
Antal K. Bejczy	
HOW FAR AWAY IS PLUG 'N' PLAY? ASSESSING THE NEAR-TERM POTENTIAL OF SONIFICATION AND AUDITORY DISPLAY	45 - 4
Robin Bargar	
THE MANY FACETS OF AUDITORY DISPLAY	71 - 5
Meera M. Blattner	
HAPTIC INTERFACES: HARDWARE, SOFTWARE, AND HUMAN PERFORMANCE	103 - 6
Mandayam A. Srinivasan	
MULTI-MODAL VIRTUAL ENVIRONMENT RESEARCH AT ARMSTRONG LABORATORY	123 - 7
Robert G. Eggleston	
THE CAVE™ AUTOMATIC VIRTUAL ENVIRONMENT: CHARACTERISTICS AND APPLICATIONS	149 - 8
Robert V. Kenyon	
VIRTUAL REALITY FOR AUTOMOTIVE DESIGN EVALUATION	169 - 9
George G. Dodd	
A SYNTHETIC DESIGN ENVIRONMENT FOR SHIP DESIGN	197 - 10
Richard R. Chipman	
OPEN VIRTUAL REALITY TESTBED ACTIVITIES	225 - 0
Sanford P. Ressler	

DISTRIBUTED AND COLLABORATIVE SYNTHETIC ENVIRONMENTS . . .	245 - 11
Chandrajit L. Bajaj and Fausto Bernardini	
INFORMATION VISUALIZATION: BEYOND TRADITIONAL ENGINEERING.	259 - 12
James J. Thomas	
THE POWER OF PEZONOMICS	295 - 3
Joel N. Orr	

Attendees

Dr. C. L. Bajaj
Department of Computer Science
Purdue University
West Lafayette, IN 47907
(317) 494-6531; Fax (317) 494-0739
E-mail: bajaj@cs.purdue.edu

Mr. Walter Baskin
Dept. E-40, Bldg. 600-1
Newport News Shipbuilding
4215 Marshall Avenue
Newport News, VA 23607
(804) 380-2662; Fax (804) 688-1073
E-mail: m.l.neumann@vivid.nns.com

Dr. Robin Bargar
Audio Dev., Virtual Environmental Group
National Center for Supercomputing Applications
Univ. of Illinois at Urbana-Champaign
5600 Beckman Institute for Advanced
Science and Technology
405 North Mathews Avenue
Urbana, IL 61801
(217) 244-4692; Fax (217) 244-2909
Email: rbargar@ncsa.uiuc.edu

Dr. Antal K. Bejczy
Mail Stop 198-219
Jet Propulsion Laboratory
4800 Oak Grove Drive
Pasadena, CA 91109
(818) 354-4568; Fax (818) 393-5007

Professor Meera Blattner
Mail Stop L540
Lawrence Livermore National Laboratory
P.O. Box 808
Livermore, CA 94551
(510) 422-3503; Fax (510) 423-4139
E-mail: blattner@ocfkms.llnl.gov

Dr. W. Scott Burton
Center for Computational Structures Technology
University of Virginia
Mail Stop 369
NASA Langley Research Center
Hampton, VA 23681
(804) 864-1992; Fax (804) 864-8089
E-mail: w.s.burton@larc.nasa.gov

Ms. Norma K. Campbell
Mail Stop 355
NASA Langley Research Center
Hampton, VA 23681
(804) 864-1131; Fax (804) 864-7722
E-mail: n.k.campbell@larc.nasa.gov

Mr. John J. Chapman
Mail Stop 423
NASA Langley Research Center
Hampton, VA 23681
(804) 864-4833; Fax (804) 864-8295
E-mail: j.j.chapman@larc.nasa.gov

Mr. Richard R. Chipman, Program Manager
Strategic Business Division
Science Applications Int. Corp.
Mail Stop 2-3-1
1710 Goodridge Drive
McLean, VA 22102
(703) 422-5670; Fax (703) 821-1134
E-mail: chipman@mclapo.saic.com

Dr. Nee-Yin Chou
CW Optics, Inc.
117 Quantico Loop
Yorktown, VA 23693-2611
(804) 867-7893; Fax (804) 868-9309
E-mail: N.CHOU@genie.gois.com

Mr. David E. Cordner
Mail Stop 157A
NASA Langley Research Center
Hampton, VA 23681
(804) 864-7325; Fax (804) 864-8342
E-mail: d.e.cordner@larc.nasa.gov

Dr. Kent T. Danielson
Center for Computational Structures Technology
University of Virginia
Mail Stop 369
NASA Langley Research Center
Hampton, VA 23681
(804) 864-8668; Fax (804) 864-8089
E-mail: k.t.danielson@larc.nasa.gov

Dr. George G. Dodd, Department Head
Analytic Process Department
General Motors Research & Development Ctr.
30500 Mound, Bldg. 1-6
P.O. Box 9055
Warren, MI 48090
(810) 986-1506; Fax (810) 986-9356
E-mail: dodd@gmr.com

Mr. George Dreybus
Dept. E-40, Bldg. 600-1
Newport News Shipbuilding
4215 Marshall Avenue
Newport News, VA 23607
(804) 380-2662; Fax (804) 688-1073
E-mail: m.l.neumann@vivid.nns.com

Mr. David Eggleston
Fakespace, Inc.
4085 Campbell Avenue
Menlo Park, CA 94025
(415) 688-1940; Fax (415) 688-1949

Dr. Robert G. Eggleston
Armstrong Laboratory
Human Engineering Division
2255 H Street
Wright Patterson Air Force Base, OH 45433
(513) 255-8764; Fax (513) 255-8752
E-mail: reggleston@falcon.al.wpafb.af.mil

Dr. Stephen R. Ellis
Mail Stop 262-2
NASA Ames Research Center
Moffett Field, CA 94035
(415) 604-6147; Fax (415) 604-3729
E-mail: Steve.Ellis@eos.arc.nasa.gov

Mr. James W. Fenbert
Mail Stop 248
NASA Langley Research Center
Hampton, VA 23681
(804) 864-5973; Fax (804) 864-3553
E-mail: J.W.Fenbert@larc.nasa.gov

Mr. F. Wallace Harrison
Mail Stop 367
NASA Langley Research Center
Hampton, VA 23681
(804) 864-6680; Fax (804) 864-4449
E-mail: f.w.harrison@larc.nasa.gov

Mr. John D. Hestenes
Division of Information, Robotics and Intelligent
Systems (IRIS)
National Science Foundation
1800 G Street, Suite 310
Washington, D.C. 20550
(703) 306-1930; Fax (703) 306-0599
(Drexel Univ., Philadelphia, PA (215) 895-
1924/2215 ; Fax (215) 895-4983
E-mail: jhestene@ccc.drexel.edu

Dr. Levent Karaoglan
Center for Computational Structures Technology
University of Virginia
Mail Stop 369
NASA Langley Research Center
Hampton, VA 23681
(804) 864-8518; Fax (804) 864-8089
E-mail: levent@puma.larc.nasa.gov

Dr. Robert V. Kenyon
Department of Electrical Engineering
and Computer Science (M/C 154)
University of Illinois
851 South Morgan, Rm 1120 SEO
Chicago, IL 60607
(312) 996-0450; Fax (312) 413-0024
E-mail: kenyon@eecs.uic.edu

Dr. Yong H. Kim
Center for Computational Structures Technology
University of Virginia
Mail Stop 369
NASA Langley Research Center
Hampton, VA 23681
(804) 864-1990; Fax (804) 864-8089
E-mail: ykim@puma.larc.nasa.gov

Dr. Doron Kishoni
College of William and Mary
34 Carriage Hill Drive
Poquoson, VA 23662
(804) 868-6570; Fax (804) 868-6057
E-mail: d.kishoni@ieee.org

Mr. Robert A. Kudlinski
Mail Stop 157
NASA Langley Research Center
Hampton, VA 23681
(804) 864-6716; Fax (804) 864-7635
E-mail: r.a.kudlinski@larc.nasa.gov

Mr. Michael C. Lightfoot
Mail Stop 237
NASA Langley Research Center
Hampton, VA 23681
(804) 864-4690; Fax (804) 864-7607
E-mail: k.severance@larc.nasa.gov

Dr. Kwanliu Ma
ICASE
Mail Stop 132C
NASA Langley Research Center
Hampton, VA 23681
(804) 864-2195; Fax (804) 864-6134
E-mail: kma@icase.edu

Dr. John B. Malone
Mail Stop 105
NASA Langley Research Center
Hampton, VA 23681
(804) 864-8983; Fax (804) 864-8915
E-mail: j.b.malone@larc.nasa.gov

Dr. Creve C. Maples
Synthetic Environment Laboratory
Mail Stop 318
Sandia National Laboratory
P.O. Box 5800
Albuquerque, NM 87185
(505) 844-2047; Fax (505) 844-2415

Mr. Michael Neumann
Dept. E-40, Bldg. 600-1
Newport News Shipbuilding
4215 Marshall Avenue
Newport News, VA 23607
(804) 380-2662; Fax (804) 688-1073
E-mail: m.l.neumann@vivid.nns.com

Dr. Ahmed K. Noor
Center for Computational Structures Technology
University of Virginia
Mail Stop 369
NASA Langley Research Center
Hampton, VA 23681
(804) 864-1978; Fax (804) 864-8089
E-mail: a.k.noor@larc.nasa.gov

Mr. Daniel A. O'Neil
George C. Marshall Space Flight Center
Mail Stop PS05
Marshall Space Flight Center, AL 35812
(205) 544-6618; Fax (205) 544-6669

Dr. Joel N. Orr, President
ORR Associates, Inc.
5224 Indian River Road, Suite 106
Virginia Beach, VA 23464
Fax (804) 495-8548
E-mail: joel_orr@mcimail.com.

Ms. Jeanne M. Peters
Center for Computational Structures Technology
University of Virginia
Mail Stop 369
NASA Langley Research Center
Hampton, VA 23681
(804) 864-1989; Fax (804) 864-8089
E-mail: j.m.peters@larc.nasa.gov

Dr. Richard J. Petti
Macsyma, Inc.
20 Academy Street
Arlington, MA 02074
(617) 646-4550; Fax (617) 646-3161
E-mail: petti@macsyma.com

Dr. Gerry D. Pollock
Center for Computational Structures Technology
University of Virginia
Mail Stop 369
NASA Langley Research Center
Hampton, VA 23681
(804) 864-2888; Fax (804) 864-8089
E-mail: gpollock@puma.larc.nasa.gov

Mr. Sanford P. Ressler
National Institute of Standards and Technology
Bldg. 225, Room B-266
Gaithersburg, MD 20899
(301) 975-3549; Fax (301) 926-3696
E-mail: sressler@nist.gov

Ms. Catherine M. Richter
Center for Computational Structures Technology
University of Virginia
Mail Stop 369
NASA Langley Research Center
Hampton, VA 23681
(804) 864-8391; Fax (804) 864-8089
E-mail: c.m.richter@larc.nasa.gov

Mr. James E. Roberts
Mail Stop 127
NASA Langley Research Center
Hampton, VA 23681
(804) 864-8483; Fax (804) 864-8467
E-mail: j.e.roberts@larc.nasa.gov

Mr. Kurt Severance
Mail Stop 125A
NASA Langley Research Center
Hampton, VA 23681
(804) 864-6715; Fax (804) 864-8910
E-mail: k.severance@larc.nasa.gov

Dr. M. A. Srinivasan
36-796 Research Laboratory of Electronics
Massachusetts Institute of Technology
50 Vassar Street
Cambridge, MA 02139
(617) 253-2512; Fax (617) 258-7003
E-mail: srinivasan@cbgrrl.mit.edu

Ms. Kathryn Stacy
Mail Stop 125
NASA Langley Research Center
Hampton, VA 23681
(804) 864-6719; Fax (804) 864-7635
E-mail: k.stacy@larc.nasa.gov

Mr. Kenneth A. Stone
Mail Stop 402
GATS
NASA Langley Research Center
Hampton, VA 23681
(804) 864-5673; Fax (804) 864-8462
E-mail: stone@hops.larc.nasa.gov

Ms. Brenda M. Sullivan
Mail Stop 463
NASA Langley Research Center
Hampton, VA 23681
(804) 864-3585; Fax (804) 864-8823
E-mail: b.m.sullivan@larc.nasa.gov

Dr. Z. Peter Szewczyk
Center for Computational Structures Technology
University of Virginia
Mail Stop 369
NASA Langley Research Center
Hampton, VA 23681
(804) 864-1991; Fax (804) 864-8089
E-mail: szez@puma.larc.nasa.gov

Mr. Geoffrey M. Tennille
Mail Stop 125
NASA Langley Research Center
Hampton, VA 23681
(804) 864-6503; Fax (804) 864-3553
E-mail: g.m.tennille@larc.nasa.gov

Dr. James J. Thomas
Technology Manager
Applied Physics Center
Battelle Pacific Northwest Laboratories
MS K7-02
P.O. Box 999
Richland, WA 99352
(509) 375-2210; Fax (509) 375-6499
E-mail: jj_thomas@ccmail.pnl.gov

Dr. Jerry H. Tucker
Mail Stop 488
NASA Langley Research Center
Hampton, VA 23681
(804) 864-7342; Fax (804) 864-7944
E-mail: J.H.Tucker@larc.nasa.gov

Mr. Samuel L. Venneri
Code XS
NASA Headquarters
Washington, D.C. 20546
(202) 358-4688; Fax (202) 358-3878

Dr. Tamer M. Wasfy
Center for Computational Structures Technology
University of Virginia
Mail Stop 369
NASA Langley Research Center
Hampton, VA 23681
(804) 864-1984; Fax (804) 864-8089
E-mail: t.wasfy@larc.nasa.gov

Dr. Brian C. Watson
Center for Computational Structures Technology
University of Virginia
Mail Stop 369
NASA Langley Research Center
Hampton, VA 23681
(804) 864-4181; Fax (804) 864-8089
E-mail: b.c.watson@larc.nasa.gov

Dr. Kang Ming Xu
Center for Computational Structures Technology
University of Virginia
Mail Stop 369
NASA Langley Research Center
Hampton, VA 23681
(804) 864-9195; Fax (804) 864-8089

INTRODUCTION

Design objectives for future high-performance engineering systems include improved manufacturability, rapid prototyping, and reduction in both qualification testing and life-cycle costs. One of the major technologies that can help in achieving these objectives is "synthetic or virtual environments." Synthetic environment is a convergence of a number of technologies including user interfaces, flight and visual simulation, and telepresence. Among the advantages of using synthetic environment facilities in design, the following three can be cited: a) bringing the users of the engineering system into the design process much earlier; b) solving multiple problems occurring during the design, at once; and c) reducing the expensive and time-consuming testing.

The joint NASA/University of Virginia workshop held in Hampton, Virginia, April 26-27, 1995 provided a forum for a wide spectrum of researchers and experts from government agencies, energy laboratories, industry and universities. An attempt was made to:

- 1) assess the state-of-technology and level of maturity of several areas in human-computer interaction, and their potential application to future high-performance engineering systems; and,
- 2) provide guidelines for focused future research leading to effective use of these facilities in the design/fabrication and operation of future high-performance engineering systems.

A list of several technologies, which collectively should comprise the next-generation design environment, was compiled from the participants. These technologies include:

- 1) routine use of immersive virtual reality environments, such as the MUSE and the CAVE facilities at the Sandia National Laboratories and the University of Illinois, and the virtual wind tunnel at NASA Ames;
- 2) selective use of sonification and haptic interfaces. Sonification facilities are suited for alert or exception notification, and for easily identified auditory icons. Haptic interfaces are useful for a wide variety of applications. This is particularly true for design problems where aesthetic considerations are paramount;
- 3) integrated product and process models (such as the "Smart Product Model" used in the ARPA-funded Simulation Based Design - SBD Project) that organize product design into a hierarchical, object-oriented structure and that bind various models, analyses and simulations together with the product data;
- 4) use of advanced "physics-based" simulations of the product's performance, manufacture, operations, etc. to evaluate the product at all stages of its development cycle and to aid in training;
- 5) extensive use of distributed, networked computing and data exchange so as to integrate geographically dispersed design and product development teams;
- 6) use of high-performance computing resources via networks on an as-needed basis;
- 7) linking of cost estimation and risk management tools with the product development tools (such as was demonstrated in the SBD Project) so as to provide the design team with continuously updated metrics to guide design trade studies;

8) high reliance on visual programming techniques to increase the productivity of the product development team;

9) advances in code integration through the use of code "wrapping" to link codes as "agents"; and

10) incorporation of DIS standards and protocols to create simulations of the product performing in concert with other elements in its intended operational environment.

The presentations at the workshop illustrated the fact that most of the aforementioned technologies are mature enough for use in product design. The challenge is to integrate them into a comprehensive design environment that promotes concurrent performance of design, engineering, manufacturing, testing and operation of a virtual prototype of the system being designed.

Since models, analyses, simulations and product/data structures vary dramatically from application to application (and from industry to industry), the next-generation design system will need to be customized to its intended application to achieve significant benefits. For each application, such as design of small spacecraft, a feasibility study can be made by conducting a product development effort to evaluate the potential benefits. The effort would use existing software (where possible), quantify both benefits and costs, and identify technology holes and necessary future research in the enabling technologies for a spacecraft design environment. Furthermore, the effort should capitalize on other, ongoing government design-system development programs and VR environments by incorporating elements and systems as they are developed.

Highlights of the Workshop

**Ahmed K. Noor
Center for Computational Structures Technology
University of Virginia
Hampton, VA**

SYNTHETIC ENVIRONMENTS (SE)

Human-computer interaction and synthetic environments have been the focus of intense efforts in recent years. A number of terms have been introduced in connection with these activities. Although there are no precise definitions of these terms, some of the principal defining ideas are indicated in Fig. 1 (see Ref. 1).

In an unmediated "normal" system the human interacts directly with the environment. In contradistinction, in a synthetic environment (CSE), a human-machine interface is used to connect the human with the physical environment. Two major categories of synthetic environments can be distinguished based on what takes place on the nonhuman side of the interface. In a teleoperator system (TS) the interface is connected to a telerobot that operates in a master-slave or supervisory control mode in a real-world environment. By contrast, in virtual environment (VE) systems, the interface is connected to a computer. In virtual reality (VR) systems, only computer simulations of the real-world environment are used. In augmented reality (AR) systems, the computer simulation is overlaid on real-world environment. The term virtual reality was coined by Jaron Lanier.

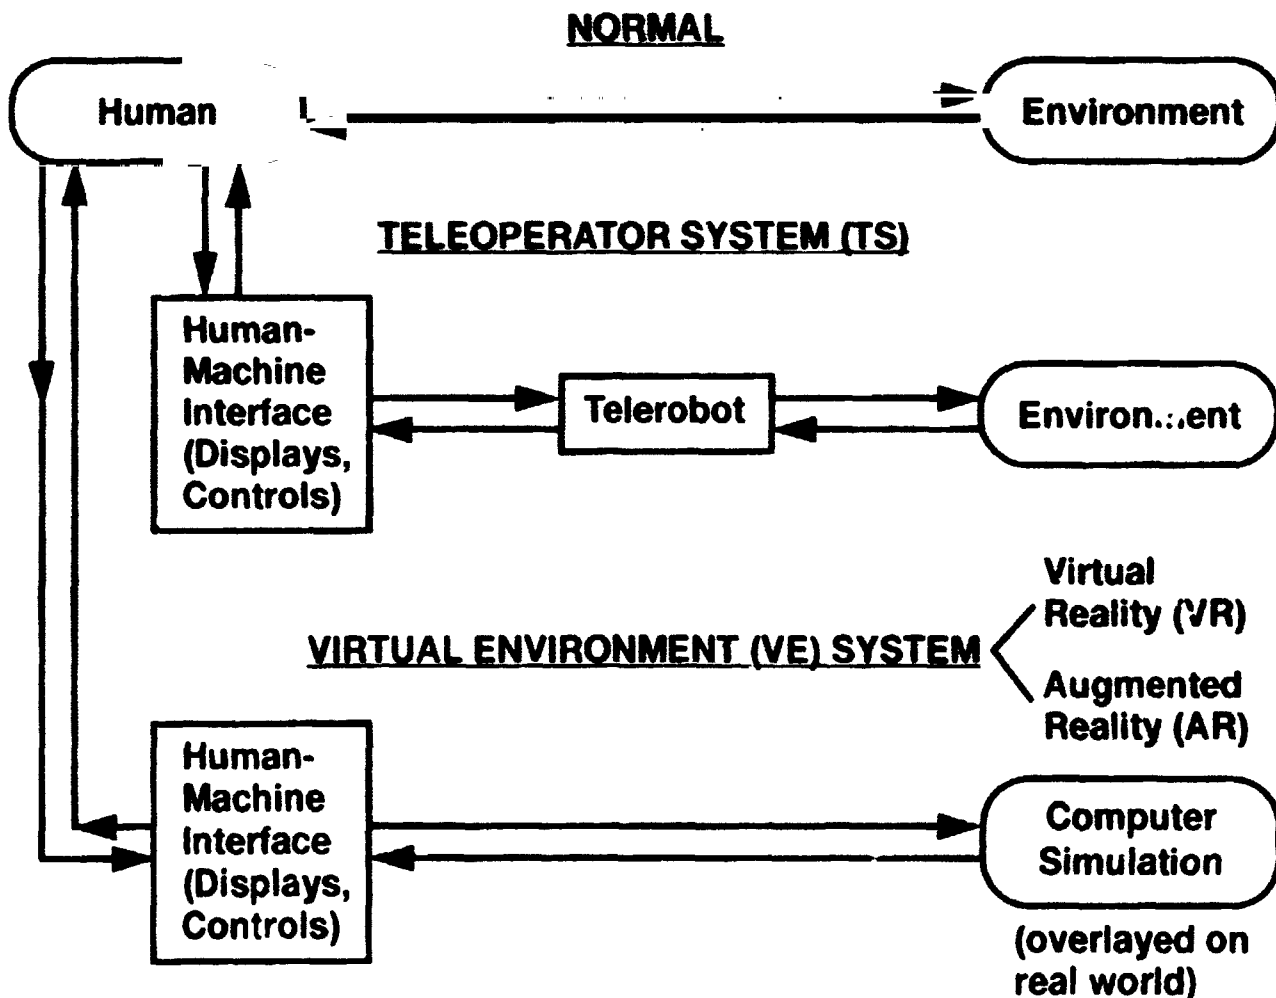


Figure 1

EVOLUTION OF NEW TECHNOLOGY

Figure 2 identifies the three phases of evolution of new technologies (such as vector and parallel computers, expert systems and synthetic environments). The first phase is that of naive euphoria, manifested by unrealistic expectations, excessive optimism, high excitement-to-accomplishment ratio, high talk-to-work ratio, and overreaction to immature technology.

The second phase is that of cynicism resulting from the frustration from not being able to meet the unrealistic expectations. After some time, the third phase is reached and the true user benefits, which *can* be achieved, are identified.

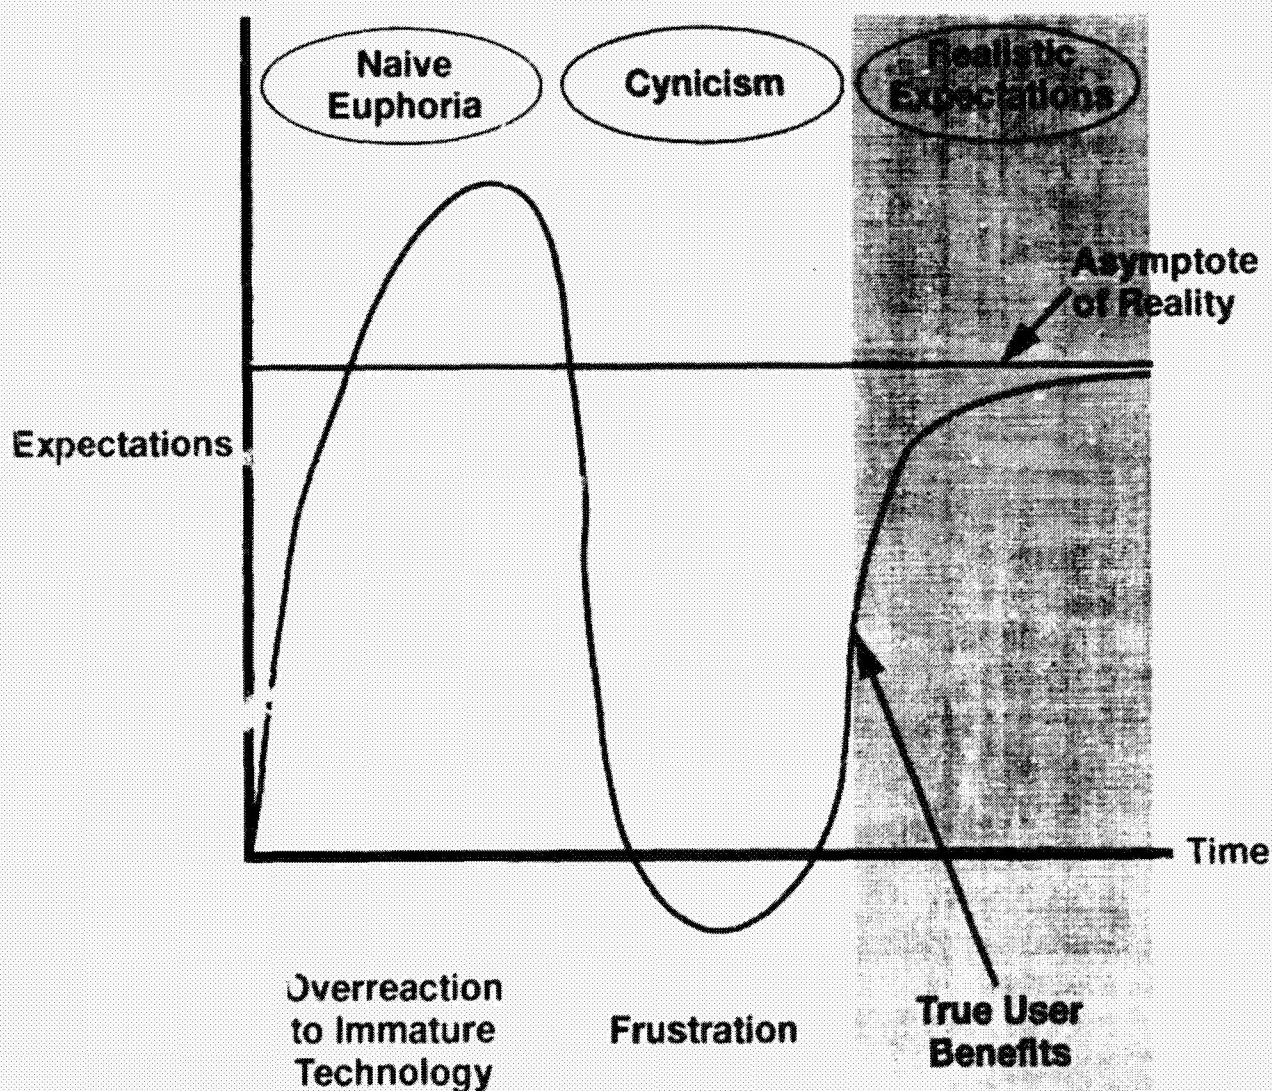


Figure 2

SOME CHARACTERISTICS OF SYNTHETIC ENVIRONMENTS

Some of the characteristics of synthetic environments are listed in Fig. 3. A major difference between animation and VR systems is the active participation of the user in the VR systems. Synthetic environment systems can make use of a wide variety of human sensing modalities and sensorimotor systems. Some of the VR systems are immersive. They create the psychological illusion of being inside the computer-generated environment through immersive display and tracking equipment (such as head-mounted display, and large screen system - such as a video projector which fills the field of view of the user). Other VR systems (e.g., desktop) use a conventional computer monitor as the output device onto which the three-dimensional environment is rendered. In some cases, stereoscopic displays are used (e.g., stereo shutter glasses). Interaction is achieved using a two-dimensional mouse, or three-dimensional devices, such as a three-dimensional mouse or a dataglove.

- **High-degree of interactivity and adaptivity**

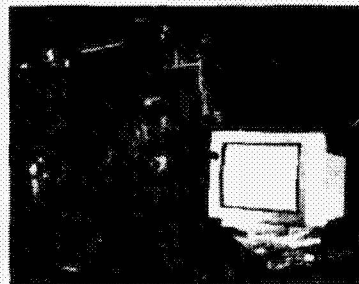
- **Multimodality (use of variety of human sensing modalities – visual, auditory displays, haptic interfaces, etc.)**



- **Immersion (strong sense of presence in the artificial environment)**



Desktop VR



Immersive VR

Figure 3

APPLICATION AREAS OF SE

Synthetic environments have high potential for a wide variety of applications. Some of these applications are listed in Fig. 4. However, with few exceptions, serious commercial applications (as opposed to research demonstrations of concept feasibility and promise) are likely to be realized in a long-term time frame (e.g., five to ten years).

In the area of education, SE systems can vastly expand the ability to access information in the larger world, to experiment, visualize, understand and interpret the information. Among the new concepts being explored are the virtual classrooms and virtual science labs. The use of SE systems for training is a natural extension of the use of simulation for training. NASA and the Department of Defense have used VE systems in training.

Visualization of complex information is another area of application of SE. An example is the virtual wind tunnel in which the user can step in and see streams of air as they pass by an electronic mockup.

In the areas of engineering design and manufacturing, the used SE can dramatically cut down the need for expensive, time-consuming physical mockups. One can design for manufacturability and maintainability. Other applications such as natural defense, telecommunications and teletravel, hazardous operations, medicine and health care, entertainment, marketing and psychology are described in Ref. 1.

- **Education and training (Virtual Classroom, Distance Learning)**
- **Handling complex information**
- **Engineering design and manufacturing**
- **National defense**
- **Telecommunication and teletravel (Network Telemeetings)**
- **Hazardous operations**
- **Medicine and health care**
- **Entertainment**
- **Marketing**
- **Psychology**

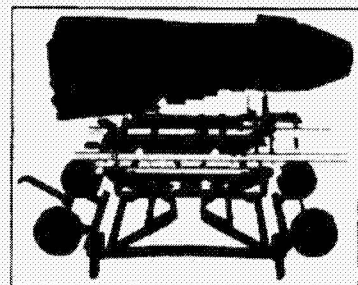
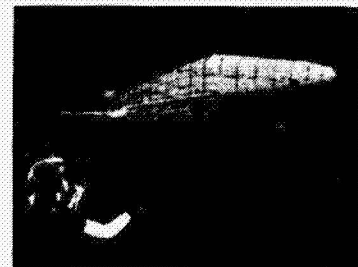
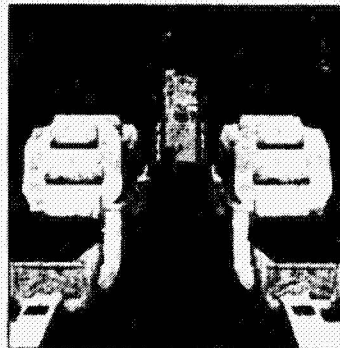


Figure 4

INFORMATION RESOURCES ON VIRTUAL REALITY

Extensive literature now exists on synthetic and virtual environments. Monographs, conference proceedings, government reports, surveys and special issues of journals have been published on the subject. In addition, a number of journals, newsletters, and short courses have been devoted to the subject. Information about virtual reality facilities and descriptions of research activities are also available on the Internet (see Figs. 5 and 6).

Books

- Introductory 26 (91-94)
- Monographs 25 (90-94), 6 (95)
- Conference Proceedings and government reports

Special issues, journal articles and surveys

- Computers and Graphics (93)
- IRIS Universe (Aug. 93)
- Computer Graphics World (Sept. 93)
- IEEE Spectrum (Oct. 93)
- Technology Review (Dec. 93)
- IEEE Computer Graphics and Applications (Jan. 94)
- Virtual Reality World - A Survey of Virtual Environments: Research in North America (93/94)

Newsletters

- Virtual Reality Review
- PC VR: Virtual Reality and the IBM Personal Computer
- CyberEdge Journal

Short Courses

Marketing Information and Catalogs

Online Resources

- Usenet
- BBS
- Lists
- WWW/ftp sites

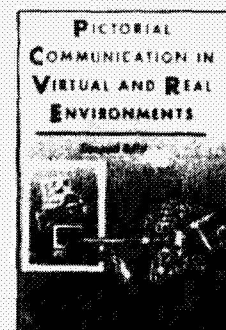
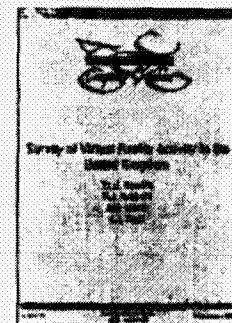
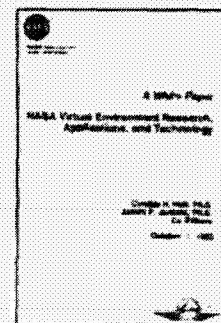
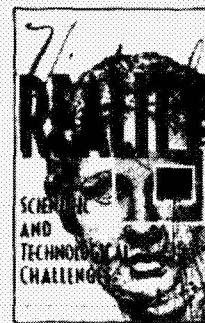


Figure 5

ASSESSMENT OF THE STATE-OF-TECHNOLOGY

Worldwide there is a significant level of activity on synthetic and virtual reality. In the U.S., the activities span over 25 universities (including Brown University, Columbia University, Georgia Institute of Technology, University of Illinois at Urbana-Champaign, Massachusetts Institute of Technology, Naval Postgraduate School, University of North Carolina at Chapel Hill, and the University of Washington), 15 federal agencies, and several industries (companies include Boeing, McDonnell Douglas, General Dynamics, General Motors and Ford).

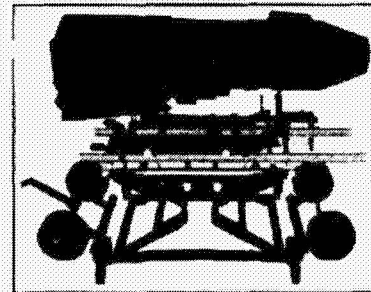
In Europe, the activities span a number of universities in the United Kingdom which are listed in Ref. 14, the Virtual Reality Center at the Fraunhofer Institute in Darmstadt, the Laboratoire d'Infographic in Lausanne, the Swedish Institute of Computer Science, the French National Institute for Audiovisual Application, and the European Space Agency in The Netherlands. Among the major European projects on virtual reality are: Elysium, Virtuosi (support for virtual organizations - virtual reality environment for office work), Telepresence, and Crimson Reality Engine.

In Japan, several academic and commercial groups are engaged in virtual reality research and development, including the University of Tokyo, NEC, Toshiba, and Matsushita.

There are an estimated 190 companies in the U.S. which produce hardware and software synthetic environment facilities.

Research activities and government programs

- U.S. (25 universities, 15 federal agencies and several companies – Boeing, McDonnell Douglas, General Dynamics, GM and Ford)
- Europe
 - Projects Elysium, Virtuosi, Telepresence, and Crimson Reality engine
- Japan
 - VR labs and medical VR programs



Current facilities and tools

- SE hardware and software facilities
- 100 large and small VR companies in the U.S

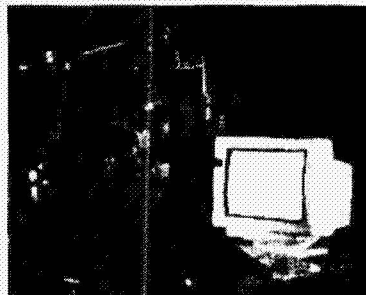


Figure 6

OBJECTIVES AND FORMAT OF THE WORKSHOP

The objectives and format of the workshop are to: a) assess the level of maturity of several human-computer interaction facilities, and their potential application to the design/fabrication and operation of future high-performance engineering systems; and b) identify future directions for research and development (see Fig. 7).

The workshop includes sixteen presentations and two panels. The presentations illuminate some of the diverse issues and provide fresh ideas for future research and development.

Objectives

- **Assess**
 - Level of maturity of several human-computer interaction facilities
 - Potential application to design/fabrication and operation of engineering systems
- **Identify future directions for research and development**

Format

- **Presentations**
- **Panels**
 - Panel 1 - Effective human-computer interfaces: dealing with complex information space
 - Panel 2 - Distributed Integrated Virtual Environments
- **Proceedings**

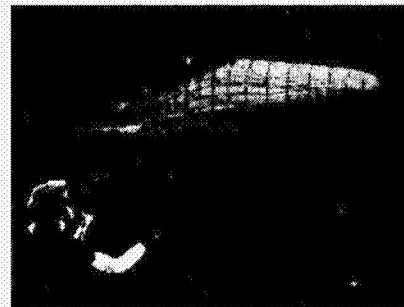


Figure 7

FUTURE DIRECTIONS FOR RESEARCH

Two of the important future research activities are listed in Fig. 8.

1. Human-computer interaction paradigms and techniques which includes the following tasks: a) mechanisms and techniques of enhancing the communication bandwidth between humans and machines; b) integrated multisensory systems; and c) cost/benefit evaluation of the impact of new paradigms on life-cycle design of high-performance engineering systems.

2. Distributed collaborative virtual environments. This includes calibration of virtual reality to physical reality; infrastructure facilities, networking and other issues; and managing of mutual awareness.

- **Human-computer interaction paradigms and techniques**

- Enhancing the communication bandwidth between humans and machines
- Integrated multisensory systems
- Impact on high-performance engineering systems (cost/benefit evaluation)

- **Distributed collaborative virtual environments**

- Calibration of VR to physical reality
- Facilities, networking and other issues
- Managing mutual awareness



Figure 8

REFERENCES

1. Durlach, N. I. and Mavor, A. S. (eds.), *Virtual Reality - Scientific and Technological Challenges*, National Academy Press, Washington, D.C., 1995.
2. Barfield, W. and Furness, T. A., III (eds.), *Virtual Environments and Advanced Interface Design*, Oxford University Press, NY, 1995.
3. Ellis, S. R. (ed.), *Pictorial Communication in Virtual and Real Environments*, Taylor & Francis Ltd, London, second edition, 1993.
4. Hamit, F., *Virtual Reality: An Exploration of Cyberspace*, Howard W. Sams/Prentice-Hall, Carmel, IN, 1993.
5. Hayward, T., *Adventures in Virtual Reality*, Que Corp., 1993.
6. Heim, M., *The Metaphysics of Virtual Reality*, Oxford University Press, NY, 1993.
7. Burdea, G. and Coiffet, P., *Virtual Reality Technology*, John Wiley and Sons, NY, 1994.
8. Larijani, L. C., *The Virtual Reality Primer*, McGraw-Hill, NY, 1993.
9. Gradecki, J., *The Virtual Reality Construction Kit*, Wiley, NY, 1994.
10. Watkins, C. D. and Marenka, S. R., *Virtual Reality Excursions: With Programs in C*, Academic Press, 1994.
11. Earnshaw, R., Jones, H. and Gigante, M. (eds.), *Virtual Reality Systems*, Academic Press, Orlando, FL, 1993.
12. Wexelblat, A., *Virtual Reality: Applications and Explorations*, Academic Press Professional, Boston, MA, 1993.
13. Singh, G., Feiner, S. K. and Thalmann, D. (eds.), *Virtual Reality Software and Technology*, Proc. VRST '94 Conference, Aug. 23-26, 1994, Singapore, World Scientific, Singapore, 1994.
14. Howard, T.L.J., Hubbold, R. J., Murta, A. D. and West, A. J., *Survey of Virtual Reality Activity in the United Kingdom*, Advisory Group on Computer Graphics, Technical Report Series, No. 27, Jan. 1995.
15. Null, C. H. and Jenkins, J. P. (eds.), *NASA Virtual Environment Research, Applications and Technology* (A White Paper), Oct. 1, 1993.
16. Gobel, M. (ed.), Special Issue on Virtual Reality, Computers and Graphics, Vol. 17, No. 6, 1993.

17. Machover, C. and Tice, S. E. (eds.), Special Issue on Virtual Reality, *IEEE Computer Graphics and Applications*, Vol. 14, No. 1, January 1994.
18. Verlinden, J. C., Kessler, D. and Hodges, L. F., *The Simple Virtual Environment (SVE) Library: User's Guide*, Technical Report GIT-GVU-93-24, Graphics, Visualization and Usability Center, Georgia Institute of Technology, Atlanta, GA, 1993.
19. Pausch, R., Crea, T. and Conway, M., "Presence, Teleoperators and Virtual Environments," *A Literature Survey for Virtual Environments: Military Flight Simulator Visual Systems and Simulator Sickness*, Vol. 1, No. 3, Summer 1992, p. 344.
20. Bryson, S. and Levit, C., "The Virtual Wind Tunnel," *IEEE Computer Graphics and Applications*, Vol. 12, No. 4, July 1992, pp. 128-137.
21. Special Issue on Virtual Reality, *IEEE Computer Graphics and Applications*, Fall 1995 (to appear).
22. *Presence, Teleoperators and Virtual Environments*, Vol. 3, No. 4, Fall 1994.

More information is available on the Internet: http://www.hitl.washington.edu/projects/knowledge_base/irvr.

**NEXT
DOCUMENT**

SI-LO
6240

THE DESIGN OF PICTORIAL INSTRUMENTS

**Stephen R. Ellis
NASA Ames Research Center
Moffett Federal Airfield, CA**

INTRODUCTION

During the last twenty years, researchers at Ames have been investigating the design of pictorial displays for civilian aerospace applications. In the mid-1980's this work sparked public imagination with the suggestion that the experience of a virtual environment provided by a head-mounted aircraft simulator could be adapted inexpensively for a wide variety of other applications. It was suggested that these head-mounted devices, called by Ivan Sutherland the "ultimate computer display," could become the ultimate computer interface. They were thought to be able to take computer users beyond the two-dimensional desktop metaphor for computer operation into a three-dimensional environment metaphor. They were billed as the "disappearing" interface, but as is currently clear, they are still far from "disappeared" (Ref. 1).

One may get a sense of the background of some of this work at Ames by very briefly reviewing its development through a few examples of studies of human performance using interactive graphical displays. The essence of this work is human interface design. The characteristic innovation is that the interface itself appears as an environment. The goal has been to use this new format to advance the safety and efficiency of human-machine interfaces in commercial aircraft and spacecraft, and such applications have materialized (Refs. 2 and 3); but other applications can extend far beyond those in aerospace.

In the course of this work the researchers at Ames have followed the long standing interest people have had in pictures as a communication medium. Pictorial representations have both symbolic and geometric components. The Lascaux Cave art, probably the oldest pictures known, emphasizes the original, probably primal interest in symbolism. But pictures also make more quantitative points. Maps are a form of pictorial representation in which the geometry is central since the message is more numerical, but we can see that even mapmakers cannot resist decorating their creations with symbols.

As the graphics capability of portable computers advanced, it became possible in the 1980's to consider pictorial display formats for electronic displays in aircraft and spacecraft. At that time investigations began of the design of aircraft traffic displays incorporating perspective projections.

The intuitive and immediate comprehensibility of such displays made them attractive for situations requiring quick interpretation; but study of them rapidly shows that effective communication requires much more than visual impact and, in particular, that the well known ambiguities of perspective images have to be counteracted (Ref. 4).

In fact, research and implementation of reference systems at Ames has determined that considerable geometric and symbolic enhancement was necessary to make a perspective-based picture useful as a display for guidance and control (Refs. 5 and 6). The aircraft display shown in Fig. 1 (lower right) is not just a pretty picture. Some of the enhancements that transform it from what one could call a spatial display into a spatial instrument, a display designed for a purpose, are described in the caption.

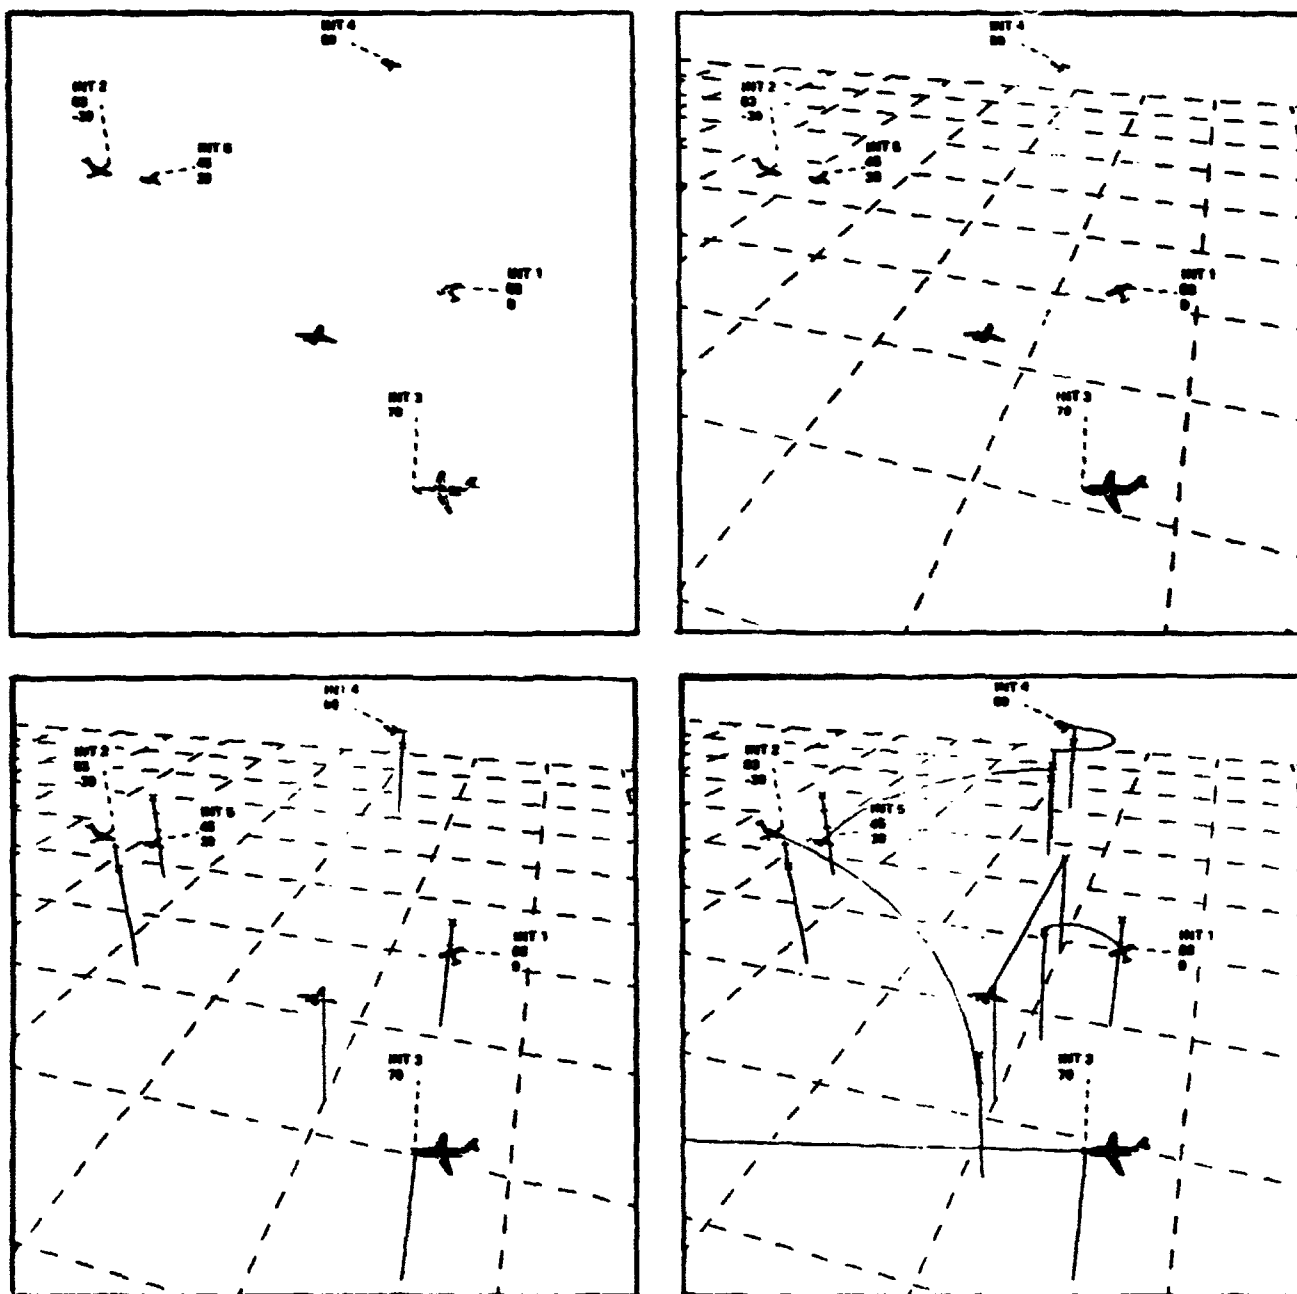


Figure 1 - Simply rendering an aircraft traffic pattern in perspective relative to a reference aircraft (center of display) does not in itself provide much spatial sense of the layout (upper left). Adding a reference grid helps but relative altitude is still hard to perceive (upper right). Adding drop-reference lines and coaltitude "x's" to show the reference aircraft's altitude against all others helps, but aircraft aspect is still ambiguous (lower left). Adding predictor lines with additional drop reference lines helps remove aspect ambiguity and resolve future traffic conflicts (lower right).
See Refs. 5 and 6 for more details.

Thus, as drawn by our E&S Picture System this display had many physically unrealizable characteristics, called geometric enhancements. These enhancements turned the display into an instrument. They allowed pilots to easily use the display to correct flight path errors, despite difficulties associated with rendering. Perspective ambiguities were only one such problem considered. But another way to avoid some of the uncertainties associated with perspective images is to jump into the picture itself.

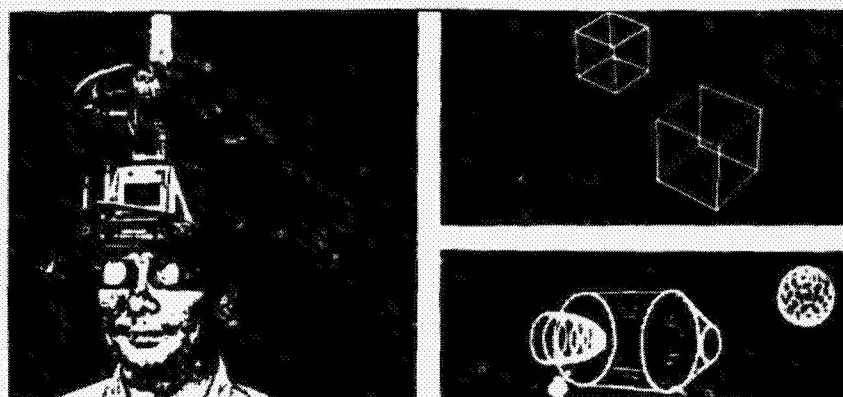
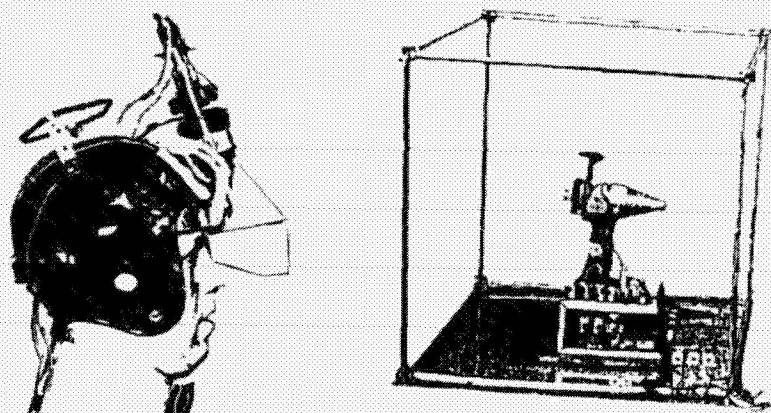


Figure 2 - Visual virtual environment display systems have three basic parts: a head-referenced visual display, head and/or body position sensors, a technique for controlling the visual display based on head and/or body movement. One of the earliest systems of this sort developed by Philco engineers (Ref. 7) used a head-mounted, binocular, virtual image viewing system, a Helmholtz coil electromagnetic head orientation sensor, and a remote TV camera slaved to head orientation to provide the visual image. Today this would be called a telepresence viewing system (upper left panel). It was used to control a remote camera (upper right panel) that was controlled by the user's head position. The first system to replace the video signal with a totally synthetic image produced through computer graphics was demonstrated by Ivan Sutherland for very simple geometric forms (lower panels) (Ref. 8).

The first implementation of a display system that allowed a viewer to do this is illustrated here in a figure from the early 1960's (Fig. 2). It was built by two Philco engineers, Brian and Comeau (Ref. 7).

This device illustrates the now familiar components of all subsequent systems using head-referenced displays. Sensors to monitor body positions, i.e., head and hand position trackers, effectors to present visual information, i.e., a miniature video display with accommodative relief, and some interlinkage hardware to connect the other two components.

These components combine in an active control loop which gives users the impression that they are located at a remote work site (Fig. 3).

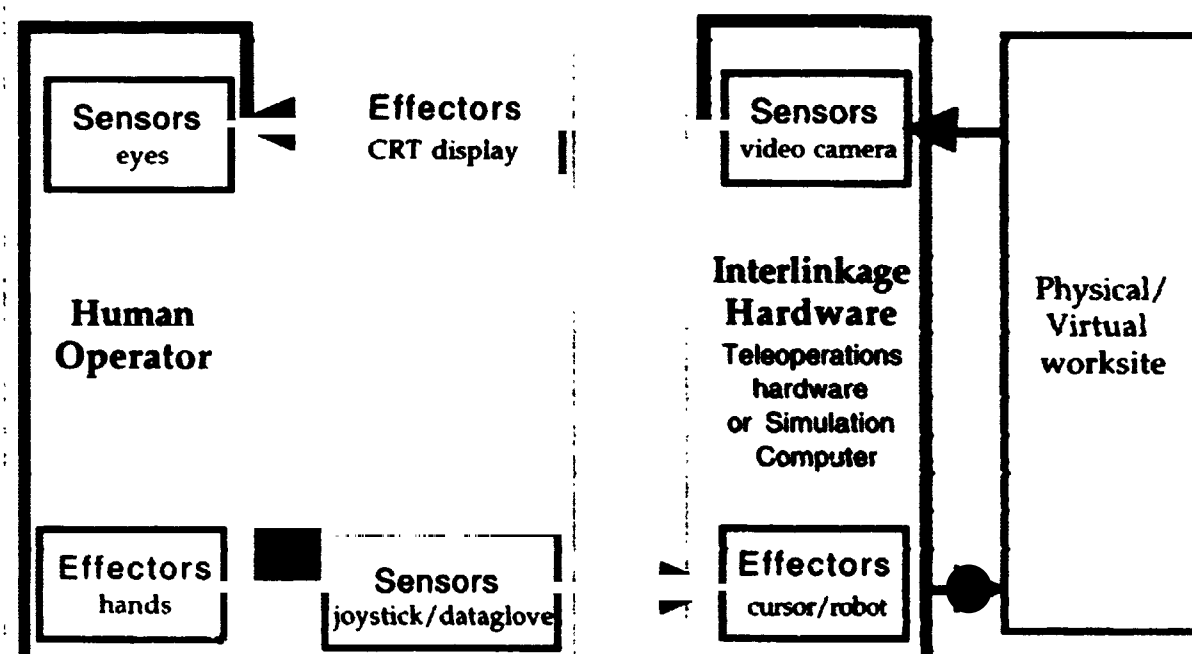


Figure 3 - Information flows in a virtual environment simulation or telepresence display.

In the case of a telepresence display the interlinkage hardware includes cameras and manipulators; this work site is real, but it could well be synthetic if the interlinkage hardware were a computer simulation.

Thus, a central aspect of head-referenced display systems is that they amount to personal simulators. Often they are simulators that are worn rather than entered. Accordingly, among the first operational display systems of this sort were indeed, head-mounted aircraft simulators (see Fig. 4).

Unlike traditional simulators associated with vehicular activity, the wearable systems are probably more uniquely suited for the simulation of manipulative activity like that associated with teleoperation. In particular, we have been interested in the study of the design of direction and distance in head-mounted displays of nearby objects; that is, objects within arms length. This application domain is relatively new since most head-mounted displays heretofore used in aerospace applications present virtual image targets at least several meters away from the user.

As an example of current design issues, we have studied the benefits of incorporating head-roll tracking for accurate presentation of direction and orientation information in a virtual environment. The need to include roll tracking is interesting since we could achieve shorter transmission and processing lags if it is left out.



Figure 4 - Though very expensive, the CAE Fiber Optic Helmet Mounted display, FOHUD (upper panel), is one of the highest-performance virtual environment systems used as a head-mounted aircraft simulator display. It can present an overall visual field $162^\circ \times 83.5^\circ$ with 5-arcmin resolution with a high resolution inset of $24^\circ \times 18^\circ$ of 1.5 arcmin resolution. It has a bright display, 20 Foot-Lambert, and a fast, optical head-tracker, 60-Hz. sampling, with accelerometer augmentation.

In our experiment we slaved a stereo camera platform having three degrees of rotational freedom to a head-mounted display. We studied the user's ability to position and orient objects in the control site to corresponding objects in the remote site (Fig. 5). Despite assertions of the benefits of matching human head kinematics through roll tracking, pitch and yaw were enough. We found that disabling roll had very little effect on performance provided subjects were not required to make very large rotations of their head with respect to their torso (Ref. 9).

In another example, we have studied perceptual issues that effect the apparent depth of nearby virtual objects presented via head-mounted displays called electronic haploscopes. They allow flexible control of the optical characteristics of binocular stimuli that may be optically superimposed in the subject's view. With this display a computer graphics image of a virtual object can be precisely calibrated so that factors that affect its apparent depth with respect to its intended depth may be studied.

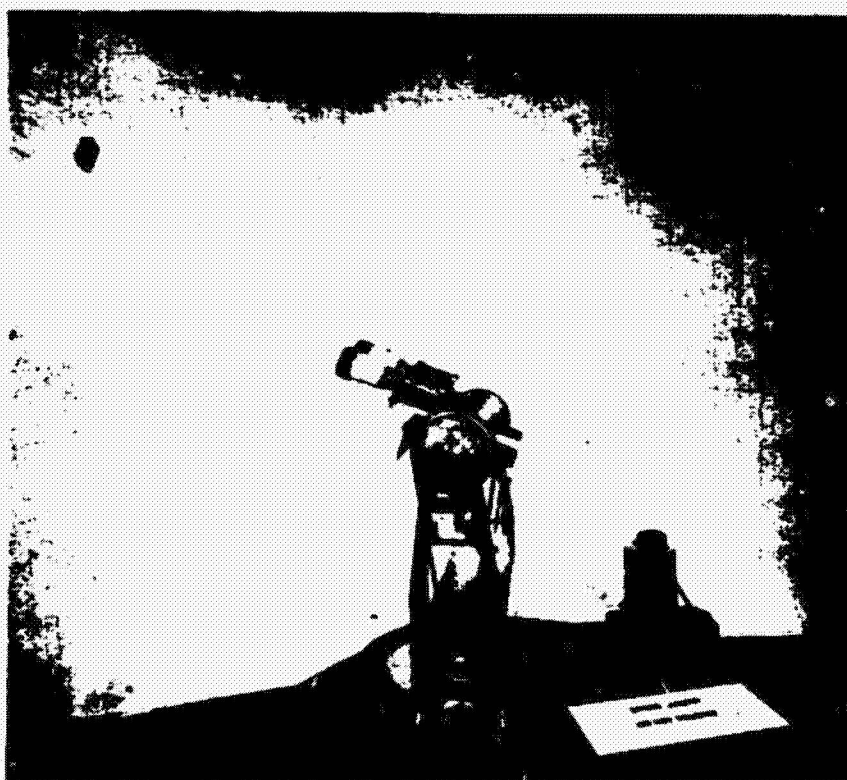


Figure 5 - The upper panel shows a three dof camera platform which is slued to the head orientation of a head-mounted display at the control site (lower panel). Studies of the improvement of the users' abilities to position and orient objects at the control site to match objects at the remote site have shown that the addition of camera roll only improves performance if large head rotations relative to the torso are forced. Pitch and yaw are usually sufficient.

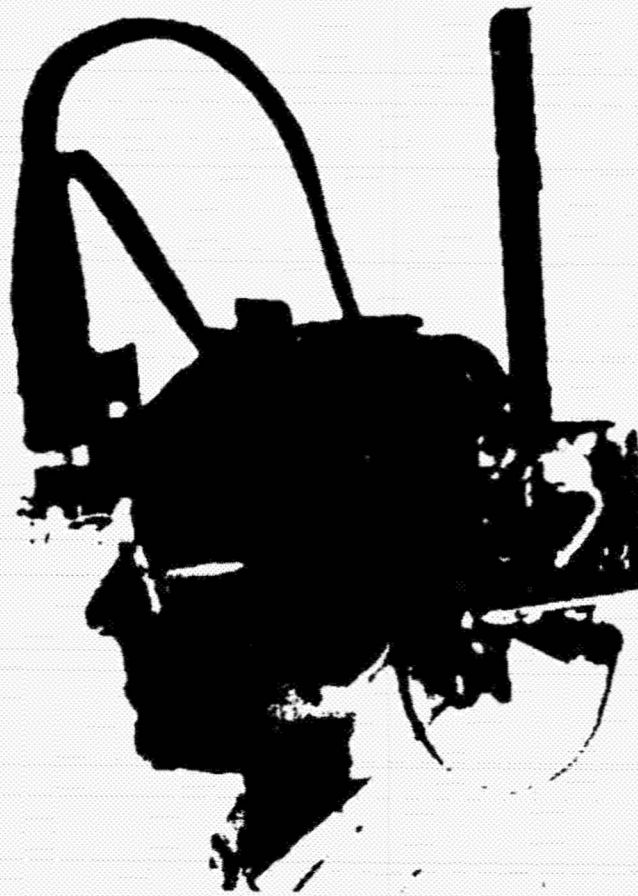


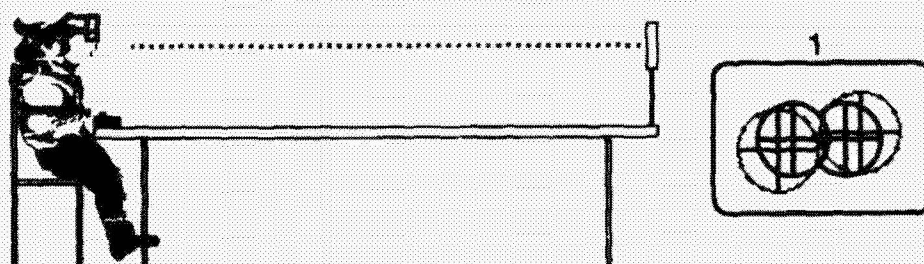
Figure 6 - An example of a head-mounted see-through display with adjustable viewing optics that allows independent variation in ocular vergence and accommodative demand making it an electronic haploscope.

Most recently, study has focused on an interesting interaction between the apparent distance to the virtual object and real surfaces over which it is superimposed. Briefly, we find that after graphics calibration and alignment, the apparent distance of the virtual object, as indicated by adjustment of a real cursor, is reduced by superposition on a real surface at the apparent distance of the object. This phenomena has clear implications for placement of virtual objects in displays for medical visualization, teleoperation and mechanical assembly. Current work is directed to describe it and to understand its cause (Ref. 10).

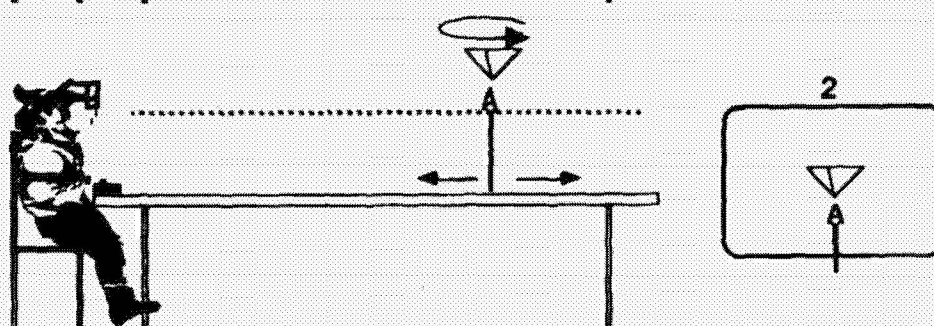
As part of interest in the dynamic aspects of the simulation loop, studies have been done on the intrinsic response latency of sensors and future studies on the role of latency, and update rates are planned for a number of perceptual and manual control phenomena, including the apparent offset of depth just mentioned.

In conclusion, it is useful to recall an observation about human interaction with pictorial displays developed from our research on the design of aerospace instruments. The phenomena most useful for the design of efficient pictorial displays are not invariably associated with the realistic rendering of the images, but rather the display and control of error. In some respects a sense of presence is not the central communication of our images. Display images, like maps, are schematic pictures with a point, often a numerical point. The principle challenge of designing them is often to determine how to portray error in a quickly interpretable manner so corrective action can be taken.

Alignment and Calibration



Cursor Adjustment to Distance of Virtual Image



Study Effects of Occlusion and Motion on Judged Position

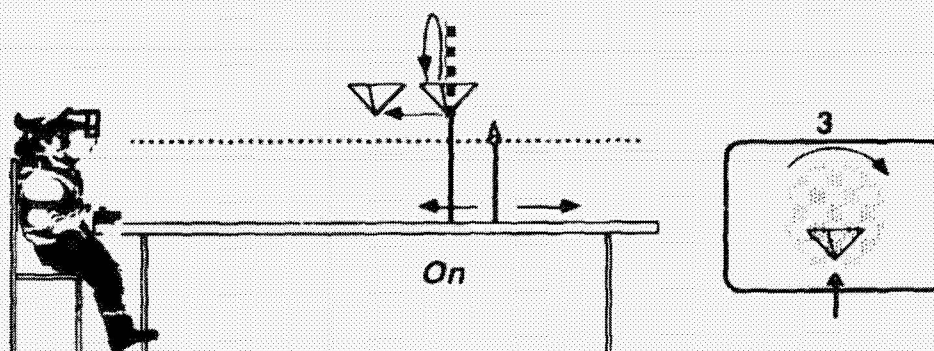


Figure 7 - an experimental scheme for examining the effects of the judged distance to computer-generated rotating virtual objects after they have been superimposed on rotating physical surfaces. After adjustment of the display parameters to correctly display visual direction for both the left and right eyes independently (top panel), the depth of a stereoscopically presented virtual object can be pointed out with a physical cursor to an accuracy of several millimeters (middle panel). Interposition of a physical surface at the judged depth of the virtual object causes the virtual object to appear to move closer to the observer (lower panel) (see Ref. 10 for experimental details).

Displays such as the one illustrated in Fig. 8, used to plan orbital maneuvers, do not depend upon photo realistic naturalness, but fluent, low-order control of animated graphical icons displaying constraints, status, and the future conditions of objects. In this case, a carefully controlled presentation of the counter-intuitive aspects of orbital mechanics was the key innovation, not making the astronauts feel that they were in orbit (Ref. 11).

This display of the counter-intuitive and nonlinear aspects of orbital maneuvering has provided a basis for the design and testing of a special format that allows *in situ* design of minimum-fuel, multi-burn maneuvers in a simulated multi-spacecraft planning environment. By focusing on the symbolic, geometric, and dynamic enhancements to synthetic environments, it provides examples of features that may be usefully incorporated into actual orbital display systems.

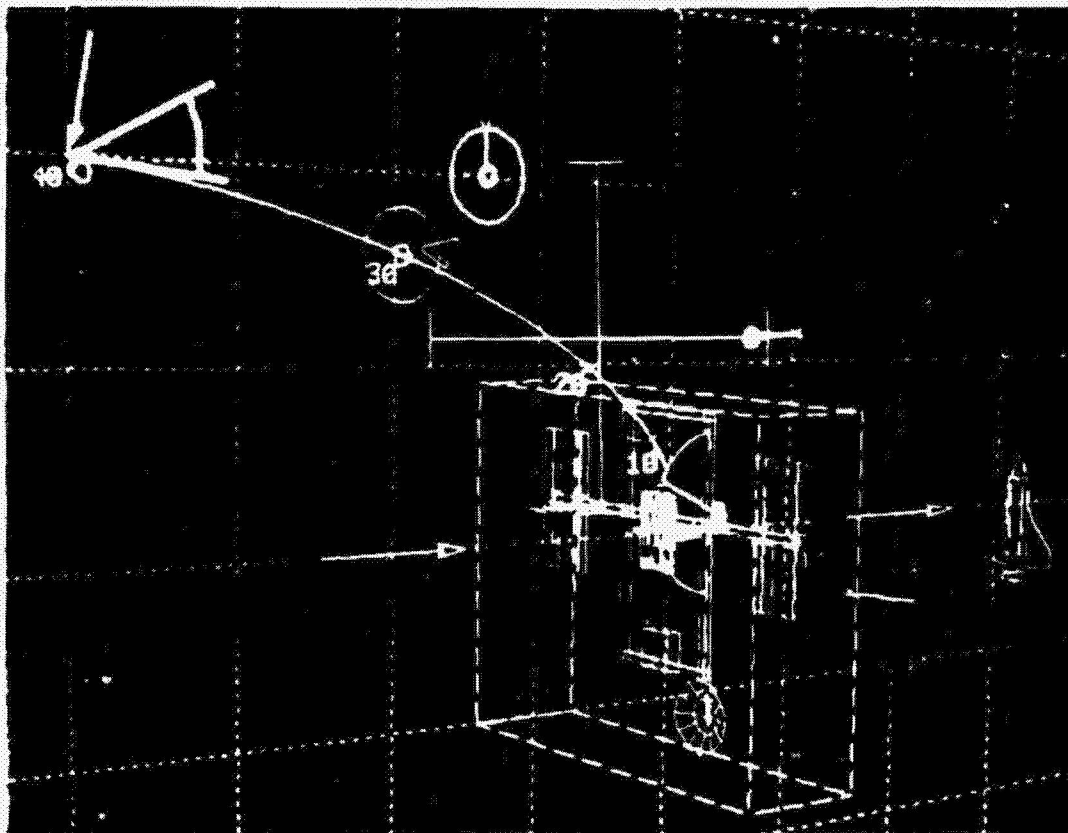


Figure 8 - A display for planning short duration orbital maneuvers subject to realistic maneuvering constraints.

The difficulty in informal planning of orbital missions arises from 1) the higher-order, nonlinear control dynamics of orbital maneuvering, 2) the counter-intuitive character of relative orbital motion, and 3) the frequent absence of stable reference points. The difficulties these characteristics pose for missions conforming to operational constraints on relative velocity, thruster plumes, and collision risk may be substantially overcome by visualizing the orbital trajectories and constraints in a pictorial, perspective display.

Though the visualization used in the display most directly assists planning by providing visual feedback to aid visualization of the trajectories and constraints, its most significant novel design features include 1) an *inverse dynamics algorithm* that reduces the order of control while also removing control nonlinearities expected from the operator and 2) a trajectory planning mode that creates, through a geometric spread-sheet, the *illusion of an inertially stable environment*. Consequently, the display is not just a "pretty picture," illustrating that computer graphics can be used to model orbital motion. It is rather a spatial instrument for interacting with orbital dynamics and now is in use in a number of laboratories around the world.

This synthetic planning environment provides the user with control of relevant static and dynamic properties of mid-course thrusts during small orbital changes allowing independent solutions to the normally coupled problems of orbital maneuvering.

This display illustrates how a synthetic environment may be defined so as to couple human problem solving abilities with the computer's computational capacities so as to enable interactive optimization of complex evaluation functions. Furthermore, it illustrates that the synthetic environments defined to enhance man-machine communication can benefit from informative symbolic, geometric, and dynamic enhancements.

As a display of error, i.e., constraint violation, this maneuvering display is distinctly non-McLuhan. This media is not the message. The focus of interest is the display and control of error, not visual impact. Thus, the paramount consideration in the design of a virtual environment as an environmental instrument is the information to be communicated and the error to be controlled.

REFERENCES

1. Ellis, Stephen R., "Nature and Origin of Virtual Environments: A Bibliographical Essay," *Computer Systems in Engineering*, Vol. 2, No. 4, 1991, pp. 321-347.
2. Brody, Adam R., Jacoby, Richard and Ellis, Stephen R., "Extravehicular Activity Self Rescue Using a Handheld Thruster," *Journal of Spacecraft and Rockets*, Vol. 29, No. 6, pp. 842-848.
3. Homan, D. J., "Virtual Reality and the Hubble Telescope," Space Programs and Technology Conference, Huntsville, AL, Sept. 27-29, 1994, AIAA Paper 94-4558.
4. Ellis, Stephen R., Kaiser, Mary K. and Grunwald, Arthur J., *Pictorial Communication in Virtual and Real Environments*, second ed., Taylor and Francis, London, 1993, pp. 28-39.
5. Ellis, Stephen R., McGreevy, Michael W., Hitchcock, Robert, "Perspective Traffic Display Format and Airline Pilot Traffic Avoidance," *Human Factors*, Vol. 29, 1987, pp. 371-382.
6. McGreevy, Michael W. and Ellis, Stephen R., *Format and Basic Geometry of a Perspective Display of Air Traffic for the Cockpit*, NASA TM-86680, Ames Research Center, Moffett Field, CA, 1991.
7. Comeau, C. P. and Brian, J. S., "Headsight Television System Provides Remote Surveillance," *Electronics*, Nov. 10, 1961, pp. 86-90.
8. Sutherland, I. E., "The Ultimate Display," *International Federation of Information Processing*, Vol. 2, 1965, p. 506.
9. Adelstein, Bernard D. and Ellis, Stephen R., "Effect of Head-Slaved Visual Image Roll on Spatial Situation Awareness," *Proceedings 37th Annual Meeting of the Human Factors and Ergonomics Society*, Santa Monica, CA, 1993, pp. 1350-1354.
10. Ellis, Stephen R., Bucher, Urs J. and Menges, Brian M., "The Relationship of Binocular Convergence and Errors in Judged Distance to Virtual Objects," *Proceedings of the International Federation of Automatic Control*, June 27-29, 1995, Boston, MA.
11. Grunwald, Arthur J. and Ellis, Stephen R., "A Visual Display Aid for Orbital Maneuvering: Design Considerations," *AIAA Journal of Guidance and Control*, Vol. 16, No. 1, pp. 139-144.

**NEXT
DOCUMENT**

82-60
C241

ALTERNATIVE DISPLAY AND INTERACTION DEVICES

M. T. Bolas, I. E. McDowall, R. X. Mead, E. R. Lorimer,
J. E. Hackbush and C. Greuel

Fakespace, Inc.
4085 Campbell Avenue
Menlo Park, CA 94025

ABSTRACT

While virtual environment systems are typically thought to consist of a head mounted display and a flex-sensing glove, alternative peripheral devices are beginning to be developed in response to application requirements. Three such alternatives are discussed: fingertip sensing gloves, fixed stereoscopic viewers, and counterbalanced head mounted displays. A subset of commercial examples that highlight each alternative is presented as well as a brief discussion of interesting engineering and implementation issues.

INTRODUCTION

For many, Virtual reality (VR) is synonymous with, and restricted to, goggles and gloves. VR was introduced to both the engineering community and the general public in this way and the image has been reinforced over time by trade journals and the press (Ref. 1). As Fig. 1 indicates, VR really is used to sell newspapers.



From Newsstands to Telephone Booth -- VR is Goggles and Gloves
Figure 1



Figure 2

This narrow definition is typical of a new field -- prejudgments are often made about a set of new technologies and ideas long before the technologies and ideas have a chance to evolve into useful configurations. The butler robot that will mix and serve drinks in the home has yet to materialize, but his non-humanoid cousins are busy building millions of cars.

This paper highlights a subset of new technologies that are not dramatically different from the familiar goggle and glove approach. Instead, these new technologies represent innovations tailored for specific application requirements. Also discussed are the advantages and limitations of each and how applications have motivated the development of these approaches.

GLOVES

The standard glove input device used for VR applications is based on the flex-sensing design made popular by the VPL Dataglove™ (shown in Fig. 3). This device measures flex along the length of each finger and thumb via a proprietary fiber-optic bend sensor (Ref. 2). A number of these flex gloves based on different technologies are commercially available. Virtual Technologies produces a high-end glove that uses 22 precision strain gauge sensors. Mattel Toys sold hundreds of thousands of flex gloves for video game use. Their glove incorporated low-cost resistive ink sensors. Exos produces a device that consists of a mechanical exoskeleton that is attached to the hand and uses rotary encoders to determine finger flex.

Typically, a graphics workstation uses the data from a flex glove to drive a kinematic model of the human hand. The computer model of the hand is rotated at the knuckles to match the corresponding flex data at each knuckle sensor of the glove. In this way a geometric representation of the overall shape of the hand is displayed.

While many VR systems incorporate such an exact graphical representation of the hand, the primary requirement for most applications is to allow the human hand to naturally interact with a virtual world. This does not always require an actual display of the hand. It does require the shape of the hand to be analyzed to recognize gestures or postures of the hand, which are then interpreted as a user command. Popular postures include making a fist to 'grab' a virtual object or pointing with the index finger to 'fly' in the direction pointed. Dynamic shape recognition goes one step further to recognize moving gestures (Ref. 3).



VPL Dataglove
Figure 3



Fakespace Pinch Glove
Figure 4

Unfortunately, flex gloves have a number of limitations when used to detect the precise hand postures used to indicate user commands. Because hands differ in shape and size, flex gloves must typically be calibrated for each user. Some flex gloves require that this calibration be repeated over time, especially if the glove slips while on the user's hand. Researchers also report that flex gloves can have 'unintentional interdependencies among the sensors' and that 'the input data of two executions of the same gesture may vary widely because of the bad repetition accuracy' (Ref. 4).

If the only task is to record a user's six degree of freedom intent, then an obvious alternative is to replace the glove with a hand-held spatially tracked switch. The Simgraphics Flying Mouse™ is a commercial example of such a system. This and similar devices work well for certain applications and can be made more effective if a representation of the interaction device is itself represented in the virtual world (Ref. 5). Another alternative is to design the VR system to function without the need for hand interaction (Ref. 6).

Nonetheless, many applications would benefit from a device that allows for a natural representation of hand interaction. Such natural interaction does not require an actual graphical representation of the hand, but it does require the computer to recognize hand actions in a way that is consistent with the user's feelings of immersion (Ref. 7).

One way to accomplish this recognition of natural gestures is to sense contact between the fingertips. Fakespace's Pinch™ Glove (shown in Fig. 4) senses the completion of a conductive path between any of the fingers and the thumb. In this fashion a number of gestures can be recognized that have natural meaning to the user. Note that flex data are not required for gesture recognition, only simple binary information between the fingertips. For example, Fig. 4 shows a 'pinching gesture' that can be used to grab a virtual object, while a 'finger snap' between the middle finger and thumb can be used to initiate an action.

A natural representation of the hand not only allows for a rapid and intuitive assimilation of a user into an environment, but there is anecdotal evidence that it makes a user feel more 'grounded' in the virtual space, thus reducing the potential for nausea and disorientation. The Pinch Glove design works with different sized hands, requires no calibration and does not drift over time.

While fingertip sensing gloves have been postulated for some time (Ref. 8), the Institute for Simulation and Training (IST) was among the first to develop such alternative glove devices. IST's first implementation was a program called Polyshop that uses two spatially tracked fingertip gloves to enable the rapid construction and manipulation of simple polygon based geometry (Ref. 9). A compelling 'tangram' puzzle demonstration using Polyshop quickly proved the utility of the interface. The Toyscout organization at IST demonstrated a number of simple games that used opposing finger gestures to fire a weapon at an attacking enemy. The speed with which young users of the game were aiming, firing and then ducking testified to the level of immersion consistent with this interface.

Fingertip sensing gloves could also be modified by:

- Adding force-sensing elements to each fingertip, thus allowing for analog force interaction with virtually grabbed objects.
- Adding flex-sensing elements to create a hybrid glove that adds hand shape representation to the fingertip sensing's quick and error-free gesture recognition.
- Using fingertip sensing elements in a hybrid glove to 'self-calibrate' flex sensors while the glove is in use.
- Adding a 'sensing work surface' to measure contact between the fingertips and the desktop, thus enabling interaction between the glove and worksurfaces such as a virtual desktop.

GOGGLES

Virtual reality systems have created a strong demand for head mounted displays (HMDs). In the past five years these displays have evolved from expensive and heavy systems primarily fashioned for military training and simulation to lightweight and inexpensive systems designed for the home computer market.

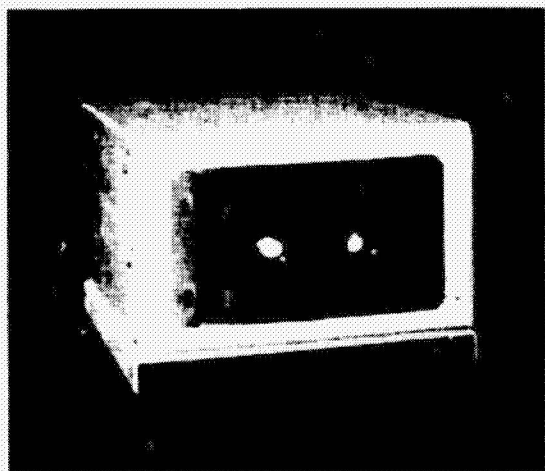
Even with the advances made, few if any HMD systems meet the following three design goals: an 80 degree field of view or greater; more than 1 million pixels per eye and weight on the user's head of less than 10 ounces. This means that for the time being, the HMD is either going to be narrow, fuzzy or heavy - characteristics that detract from an immersive experience.

While many applications require the use of an HMD system despite their current drawbacks, not all do. For these other applications, alternative viewing systems can meet the three basic design goals outlined above while also providing an additional set of benefits. Interestingly, the two alternative systems presented here are based on configurations that

existed before the modern HMD. These are the fixed stereoscopic viewer and the counterbalanced head mounted display.

Fixed Stereoscopic Viewers

Similar in many ways to the early nickelodeon display, an effective alternative to HMDs can be found in the fixed stereoscopic viewer (FSV).



Fakespace Immersive Stereo Viewer

Figure 5



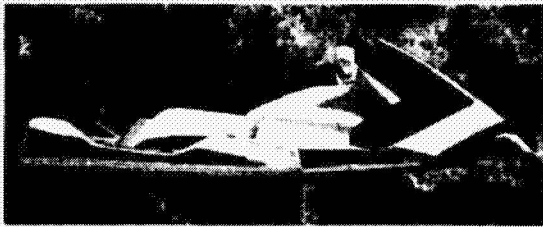
Telepresence Virtual Brewery Exhibit

Figure 6

Because of the form factor, an FSV can meet the design goals outlined above. The Fakespace Immersive Stereo Viewer (ISV™) (shown in Fig. 5) offers over 1.2 million pixels per eye with a selectable field of view ranging from 30 to 100 degrees. Fixed displays of this type are naturally robust, economical and provide high user throughput. These qualities make the FSV ideal for a location based entertainment venue or other such public installation.

Figure 6 is an image from the Virtual Brewery project installed at the Sapporo Headquarters in Japan by Telepresence Research. This installation is comprised of twelve Fakespace Immersive Stereo Viewer systems coupled with a Fakespace BOOM3C™ display, an SGI Onyx™ computer and a Crystal River Acoustetron II™ (Ref. 10). While one patron uses the BOOM3C to navigate and control the point of view (POV) through a virtual brewery, twelve additional patrons can 'go along for the ride' by looking into the Immersive Stereo Viewers. This arrangement allows for over 1000 people a day to experience the virtual environment in an easy and comfortable way.

It is interesting to note that even though the user cannot use head motion to control the POV, the experience obtained with a fixed stereoscopic viewer can be quite immersive. It retains the feeling of a first person POV except that viewpoint changes are attributed to an overall shift of the frame of reference - almost like being a passenger in a vehicle. This is critical as it allows the user to feel grounded as opposed to nauseous.



Greystone Mercury Platform
Figure 7



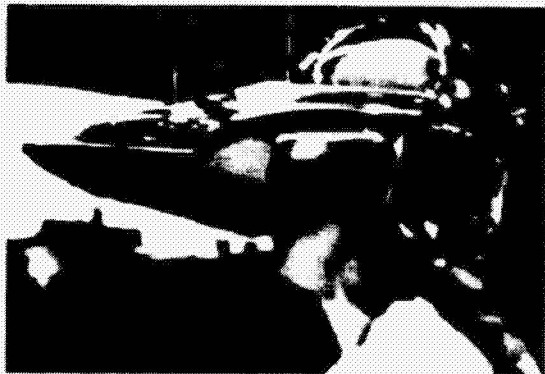
Nintendo Virtual Boy
Figure 8

Fixed stereoscopic viewers need not remain completely fixed. Greystone Technology recently released the Mercury Platform™. This system couples a Fakespace ISV display with an air powered three degree-of-freedom motion base and a modified motorcycle grip for navigation. Users of the system sit down in a fashion similar to a motorcycle and peer into the ISV display. As the user navigates through the virtual environment, the platform tilts and rolls to simulate motions corresponding to the action in the virtual environment. The fixed nature of the ISV works perfectly with the design of the overall system and takes advantage of the vehicle-based immersive feelings described above.

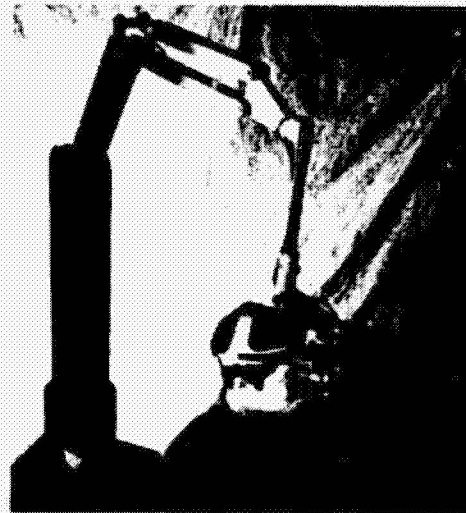
Perhaps the largest provider of FSVs will soon be the Nintendo Corporation. Nintendo has finished work on the Virtual Boy™ system that uses two monochromatic displays mounted in a table-top viewing system as seen in Fig. 8. The system was shown at the 1995 Winter Consumer Electronics Show and allows the user to control the immersive view via a hand-held controller while peering into the table-top stereoscopic display. Users have commented on the sharp images and high quality stereoscopic effect - qualities that the fixed display form factor easily achieves.

Counterbalanced Head Mounted Displays

Counterbalanced displays that are head mounted, as opposed to head coupled (Ref. 11), is the last alternative technology considered here. Counterbalanced head mounted displays have been used since the early days of computer simulation to achieve high visual and tracking performance without the limitation of weight constraints.



Disney Imagineering Aladdin Adventure
Figure 9



Fakespace FS2 Simulation System
Figure 10

Disney Imagineering has installed a CRT based stereoscopic display system at Epcot Center that is counterbalanced via an airspring and a pair of supporting cables (shown in Fig. 9). These cables ensure that the display cannot be dropped to the floor. The added weight made possible by counterbalancing allows for a rugged carbon-fiber shell that protects against public abuse. A detachable head strap increases the overall system throughput. Magnetic tracking is used as opposed to mechanical tracking due to the flexible nature of the cable counterbalance.

Fakespace produces the FS2™ Simulation System (shown in Fig. 10). The FS2 is a full six degree of freedom counterbalanced display and tracking system. The system uses lightweight materials and three independent spring counterbalances to minimize inertia and eliminate head supported weight. Noiseless mechanical tracking is achieved via optical shaft encoders while user-selectable optics achieve over a 100 degree field of view. The display attaches to the user via a head strap. The FS2 is optimized to be a seated simulation system with the automotive and simulation industries as early adopters. The system was designed to work well within the tight confines of an automobile interior. In particular, the head and structure are shaped to reduce potential interference problems with a vehicle's steering wheel, headrest and seat.

Counterbalanced head mounted displays are extremely comfortable because they exert no weight on the user's head - effectively they weigh zero ounces. This is important for many simulation systems where the operators may be immersed for hours at a time. Because of inertial effects, however, the mass of the displays and the distribution of the mass cannot be ignored. Both the FS2 and the Disney displays were designed to minimize mass without sacrificing display quality. As a result, the inertial effects are roughly the same as those for commercially available CRT based head mounted displays while maintaining fields of view of over 80 degrees with resolutions of more than 1.2 million pixels per eye.

CONCLUSION

Virtual reality has opened up a powerful and effective set of new techniques and technologies for delivering immersive experiences. As the science and art of designing these experiences matures, it becomes clear that the engineer and designer should look toward the constraints imposed by specific applications for guidance and insight. In many cases, the shedding of prior technical assumptions leads to tools that surpass expectations. The field of virtual reality is beginning to flourish with exciting and tantalizing immersive experiences and tools.

ACKNOWLEDGMENTS

The authors wish to thank Ken Socks of the Chrysler Corporation, Scott Fisher of Telepresence Research, and Bernie Crow of Greystone Technologies for their support and design insights.

REFERENCES

1. Greer, J., "Helmut of Dreams," *San Jose Mercury News*, Jan. 12, 1986, p. 3A.
2. Wise, S., Gardner, W., Sabelman, E., Valainis, E., Wong, Y., Glass, K., Drace, J. and Rosen, J. M., "Evaluation of a Fiber Optic Glove for Semi-automated Goniometric Measurements," *Journal of Rehabilitation Research and Development*, Vol. 27, No. 4, 1990, pp. 411-424.
3. Jacoby, R., Ferneau, M. and Humphries, J., "Gestural Interaction in a Virtual Environment," *Stereoscopic Displays and Virtual Reality Systems*, S. S. Fisher, J. O. Merritt and M. T. Bolas (eds.), Proc. SPIE, Vol. 2177, 1994, pp. 355-364.
4. Boehm, K., Broll, W. and Sokolewicz, M., "Dynamic Gesture Recognition Using Neural Networks: A Fundament for Advanced Interaction Construction," *Stereoscopic Displays and Virtual Reality Systems*, S. S. Fisher, J. O. Merritt and M. T. Bolas (eds.), Proc. SPIE, Vol. 2177, 1994, pp. 336-346.
5. Sowizral, H. A., "Interacting with Virtual Environments Using Augmented Virtual Tools," *Stereoscopic Displays and Virtual Reality Systems*, S. S. Fisher, J. O. Merritt and M. T. Bolas (eds.), Proc. SPIE, Vol. 2177, 1994, pp. 409-416.
6. Darken, R., "Hands-off Interaction with Menus in Virtual Spaces," *Stereoscopic Displays and Virtual Reality Systems*, S. S. Fisher, J. O. Merritt and M. T. Bolas (eds.), Proc. SPIE, Vol. 2177, 1994, pp. 365-371.
7. Bolas, M. T., *Design and Virtual Environments*, M.S. Thesis, Design Division, Stanford University, 1989.

8. Stone, P. Personal Conversation, VIEW Lab System, Human Factors Division, NASA Ames Research Center, Moffett Field, CA, 1988.
9. Mapes, D., *Two-Handed Virtual Environment Interface, Improving Productivity for Object Manipulation and Viewpoint Movement in Three-Dimensional Worlds*, M.S. Thesis, Computer Science Department, University of Central Florida, 1994.
10. Buckert-Donelson, A., "Scott Fisher," *VR World*, Vol. 3, No. 2, 1995, pp. 28-32.
11. McDowall, I. E., Bolas, M., Pieper, S., Fisher, S. S. and Humphries, J., "Implementation and Integration of a Counterbalanced CRT-Based Stereoscopic Display for Interactive Viewpoint Control in Virtual Environment Applications," *Stereoscopic Displays and Virtual Reality Systems*, S. S. Fisher and J. O. Merritt (eds.), Proc. SPIE, Vol. 1256, 1990.

**NEXT
DOCUMENT**

15
6242

COMPUTATIONAL VIRTUAL REALITY (VR) AS A HUMAN-COMPUTER INTERFACE IN THE OPERATION OF TELEROBOTIC SYSTEMS

**Antal K. Bejczy
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109**

This presentation focuses on the application of computer graphics or "virtual reality" techniques as a human-computer interface tool in the operation of telerobotic systems. VR techniques offer very valuable task visualization aids for planning, previewing and predicting robotic actions, operator training, and for visual perception of non-visible events like contact forces in robotic tasks. The utility of computer graphics in telerobotic operation can be significantly enhanced by high-fidelity calibration of virtual reality images to actual TV camera images. This calibration will even permit the creation of artificial (synthetic) views of task scenes for which no TV camera views are available.

OVERVIEW

- **VR as real-time control technique; fusion/calibration issues**
- **VR as training method; kinesthetic/tactile features/interfaces**
- **VR as task logistician; protocols and audio/vocal interfaces**
- **VR as visualization tool for abstract, non-visible things; displays**
- **Conclusion**

VR AS REAL-TIME CONTROL TECHNIQUE WHEN FUSED WITH TV CAMERA IMAGES

Task visualization is a key problem in teleoperation since most of the operator's control decisions are based on visual information. The capability of previewing motions enhances the quality of teleoperation by reducing trial-and-error approaches in the hardware control and by increasing the operator's confidence in control decision making during task execution. Predicting the consequences of motion commands under communication time delay permits the command of longer and safer action segments as opposed to the command of short action segments adopted in the move-and-wait control strategy in time-delayed teleoperation without predictive displays.

Fusion of graphics and TV camera images can be generated by overlaying graphics images over actual TV camera images. A high-fidelity overlay requires a high-fidelity TV camera calibration and object localization. For this purpose, a reliable operator-interactive camera calibration and object localization technique has been developed at JPL during the past few years. It currently uses a point-to-point mapping procedure, and the computation of the camera calibration parameters is based on the ideal pinhole model of image formation by the camera. The technique uses the robot arm as the calibration fixture and assumes the use of a few selectable, good static TV camera views. The technique was demonstrated to a broad audience in May 1993 when a JPL control station was connected to an ORU exchange mock-up 4000 km away at GSFC through the NASA-select TV channel and the Internet computer network. The task was successfully performed under varying communication time delay conditions.

The calibration technique and its application potential gained technical acceptance and is now being commercialized through a technical transfer agreement with DENEb Robotics, Inc. More on this calibration technique and its demonstration in Refs. 1 and 2. A narrated VCR tape is available on the demonstration in the JPL Audio-Visual Library (no. AVC-93-165C1D).

- **Motivation**
 - **Communication time delay: predict actions**
 - **Planning complex tasks: preview actions**
 - **Intelligent automation: supervise/monitor actions**
- **Method of fusing VR images with TV images: create high-fidelity overlays of graphics images over TV camera images**
 - **Calibration technique: point or feature mapping**
 - **Motion control of graphics overlays**
- **Extra benefit: enables artificial or synthetic views or scenes**

(See VCR tape for performance results)

VR AS TRAINING METHOD WITH KINESTHETIC/TACTILE FEATURES/INTERFACES

Operator training using a VR display system is a convenient tool for initial familiarization of the operator with the teleoperated system without actually turning the hardware system on. Using proper physical modeling, even sensors and sensor fusion can be simulated and graphically shown to the operator. Computer graphics simulation of proximity sensor signals is a relatively simple task since it only implies the computation of distance from a fixed point of a moving robot hand in a given (computed) direction to the nearest environment surface in the graphics "world model". Force-torque sensor signals can also be simulated by computing virtual contact forces and torques for given geometric contact models using a spring or a spring plus damper description of the actual contact interaction. For some detail, see Ref. 3.

Modeling of soft things, like tissues, and graphically showing their deformations as a function of pressure is a very demanding undertaking. Some useful information on this topic can be found in Ref. 4.

- **Modeling of contact forces/moments and tactile area pressure - and showing them graphically**
- **Modeling of "soft" things, like tissues - and showing their deformations graphically as a function of pressure**
- **Application potential of manual force/moment and tactile feedback from VR interaction scenes is increasing for**
 - **training operators (and surgeons)**
 - **"sensitive" teleprogramming**

VR AS VISUALIZATION TOOL FOR DISPLAYING ABSTRACT, NON-VISIBLE THINGS

Visualization of non-visible events enables a graphical representation of different non-visual sensor data and helps management of complex systems by providing a suitable graphical description of a multi-dimensional system state. For instance, the constrained and orientation restricted motion space of a dual-arm robot working in a closed kinematic chain configuration can be visualized as a complex 3-D object with hidden unreachable holes or cavities of varying shapes. An automated visualization method has been developed to find and visually represent this complex geometric object from a computed numerical data base. The method is an inverse computer vision technique in the sense that it creates rather than recognizes visual forms. More on this can be found in Ref. 5.

- Several abstract, non-visible things can be present during telerobotic operations originating from
 - Internal system constraints - which can translate to complex task space constraints
 - Different multidimensional sensor data spaces - which, by some mapping, contain task/subtask goals as single event points or restricted small volumes
- Visualization of non-visible things may require the use of
 - Complex computational procedures
 - Some artistic creativity for designing graphic forms
- An example: visualization of a constrained dual-arm geometric work space treated as an inverse computer vision problem

VR AS TASK LOGISTICIAN SURROGATE WITH AUDIO/VOCAL INTERFACES

In an emerging field of R&D, researchers embed task protocols or task scripts within a suitable VR representation. This requires to map a sequence of VR scenes to a sequence of required actions. In a related effort, researchers already initiated the idea of providing the operator with performance feedback messages on the operator interface graphics, derived from a stored model of the task execution protocol. A key element of such advanced feedback tool to the operator is a program that can follow the evolution of a teleoperated task by segmenting the sensory data stream into appropriate task performance phases. Task segmentation programs have already been implemented using Hidden Markov Model representations (Ref. 6) and Neural Network Architecture (Ref. 7) with very promising results.

- VR with embedded task protocols/scripts - emerging field of R&D work
 - Mapping the sequence of VR scenes to sequence of required actions
 - Mapping the actual performance shown on VR scenes to required performance
- Audio/vocal interface to VR-embedded task logistician is possible and desirable, with some operational restrictions
 - Fixed content both ways
 - Short statements both ways
- Consideration of human factors

CONCLUSION

The application of Virtual Reality techniques/tools in the operation of telerobotic systems enables the performance of more tasks, safer, faster, and inherently cheaper.

REFERENCES

1. Kim, W. S., "Virtual Reality Calibration for Telerobotic Servicing," *Proceedings of the IEEE International Conference on Robotics and Automation*, San Diego, CA, May 1994.
2. Kim, W. S. and Bejczy, A. K., "Demonstration of a High-Fidelity Predictive/Preview Display Technique for Telerobotic Servicing in Space," *IEEE Transactions on Robotics and Automation*, Oct., 1993.
3. Kim, W. S. and Bejczy, A. K., "Graphics Displays for Operator Aid in Telemanipulation," *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, Charlottesville, VA, Oct. 1991.
4. *IEEE Computer Graphics and Applications Magazine*, a bimonthly publication, issues of 1992 through 1994.
5. Tarn, T. J., De, P. K., Bejczy, A. K. and Li, Z., "Visualization of Dual-Arm Robot Motion Space Under Kinematic Constraints," *International Journal of Robotics Research*, MIT Press, Vol. 14, No. 1, Feb. 1995.
6. Hannaford, B. and Lee, P., "Hidden Markov Model Analysis of Force-Torque Information in Telemanipulation," *International Journal of Robotics Research*, MIT Press, Vol. 10, No. 5, Oct. 1991.
7. Fiorini, P., et al, "Neural Networks for Segmentation of Teleoperation Tasks," *PRESENCE: Teleoperators and Virtual Environments*, MIT Press, Vol. 2, No. 1, Winter 1993.

**NEXT
DOCUMENT**

94 60
6245

HOW FAR AWAY IS PLUG 'N' PLAY? **Assessing the Near-Term Potential of Sonification and Auditory Display**

Robin Bargar
Virtual Environment Group
National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign

INTRODUCTION

Sound is gradually making its way into virtual environments (VE). This presentation addresses the state of the sonic arts in scientific computing and VE, analyzes research challenges facing sound computation, and offers suggestions regarding tools we might expect to become available during the next few years. Sound immerses us in an acoustic world of rhythmic and melodic messages and environmental and spatial cues. Remove the sounds in our real world and we will be less certain where we are. When sound is included in VE, users begin to rely upon it for similar environmental orientation.

Since VE's are predominantly graphical display environments, we include discussion of sound relative to the computation and display of visual information. For many of us, the cinema provides formative experiences of sound in visual environments. The cinematic model creates strong expectations regarding the roles sound plays and the places we will be able to hear it. Sounds in VE fill many cinematic roles, giving an environment a more continuous sense of presence and providing information to enhance or reinforce visual display.

A list of classes of audio functionality in VE includes *sonification* - the use of sound to represent data from numerical models; 3D auditory display (*spatialization* and *localization*, also called *externalization*); navigation cues for positional orientation and for finding items or regions inside large spaces; voice recognition for controlling the computer; external communications between users in different spaces; and feedback to the user concerning his own actions or the state of the application interface.

To effectively convey this considerable variety of signals, we apply principles of acoustic design to ensure the messages are neither confusing nor competing. Acoustic design requires the talents of musicians and composers to ensure a listener does not experience auditory fatigue. At NCSA we approach the design of auditory experience through a comprehensive structure for messages, and message interplay we refer to as an *Automated Sound Environment*. We implement classes of auditory messages as high-level functions in a software environment for rendering sounds. Our research addresses four engineering and communication challenges: real-time sound synthesis, real-time signal processing and localization, interactive control of high-dimensional systems, and synchronization of sound and graphics. Each of these represents a set of hardware-software engines needed by the general VE community in order to effectively use sound. Such engines are not at this time commercially available. In the following pages we discuss some of the principles involved in these tools, practical issues surrounding their implementation, and examples of their application in working VE systems.

AN ACOUSTIC OBSERVATION-FEEDBACK CYCLE

Observation in VE depends upon interaction between an observer and a computational model. Numerical computation is reflected upon through the cognitive process of an observer. To achieve a reflection we need an auditory display interface and a control input from the observer to the computation. Observation includes the control gestures input from an observer investigating the system and the cognitive processing of acoustic feedback generated by the resulting state of the computational model. "Sounds" and "events" indicate the acoustic signals from the interface have been transformed into auditory signals by a listener. The terms *qualitative* and *quantitative* denote this transformation. In distinction to a practice of referring to the "qualities" of numerical data, our proposition is that numerical models have intrinsic properties, but these properties do not have "qualities" until they are perceived through the actions of an observer and an interface [1].

Bringing numerical data into cognition

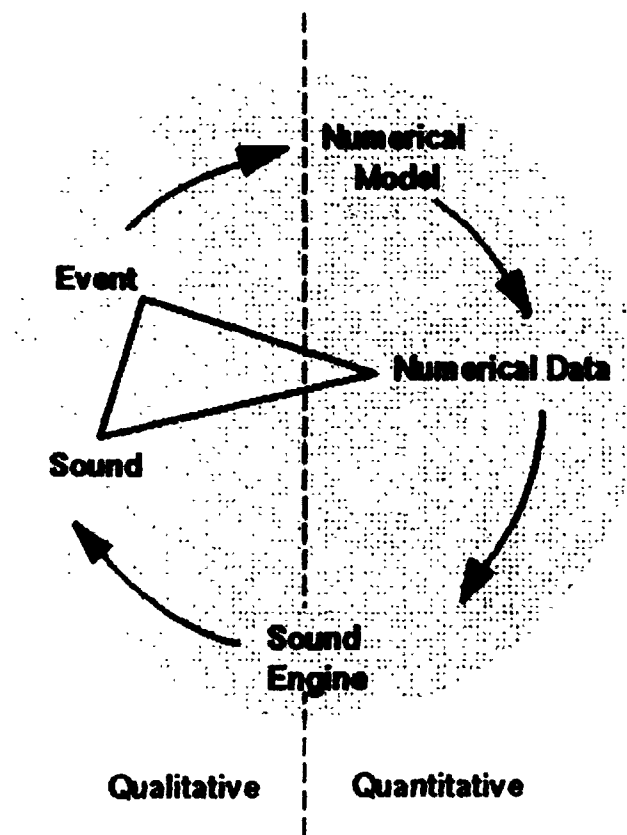


Figure 1

CONTROL FLOW ARCHITECTURE

Sonification is the rendering in sound of scientific data from a numerical model. It is one of the least explored functions of auditory display. Computer-synthesized sound is controlled by numerical data so it is possible to construct a control flow from a scientific model to a sound synthesis engine. However, there is no guarantee the scientific data will produce intelligible auditory information. A sound designer determines an appropriate mapping between the two systems. The diagram below accounts for two design stages: (1) the creation of a sound synthesis engine capable of producing a known and controllable range of sounds, and (2) the creation of an expressive relationship between the sound synthesis capability and the characteristics of the scientific data.

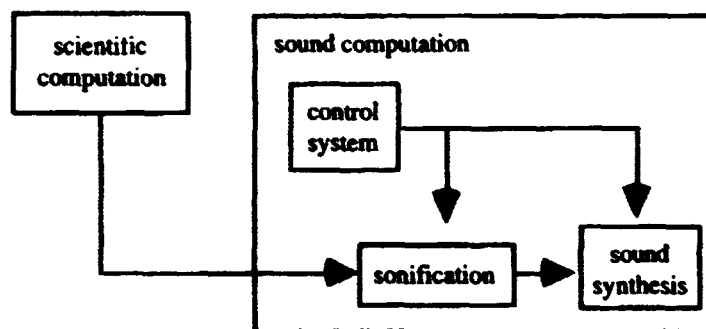


Figure 2

WHAT WE HEAR IN AN ACOUSTIC SIGNAL

Sound is characterized by energy distribution in the frequency domain and by rapid changes in the time domain. Sampling rates of 48 kHz or better are needed to encode a signal compatible to the perceptual range of the human ear. To observe the structure of sound we can decompose a signal into a series of discrete short-time Fourier transforms. Our ears are remarkably sensitive to small changes in energy in the frequency domain, $c - time$. The diagram below shows the structure of a small sample of the steady-state portion of a tone played by a trumpet. The physical structure of the trumpet provides a resonating column of air; its resonant characteristics can be seen in the regular distribution of energy peaks in the frequency domain. These energy peaks are called *partials* or *harmonics*. They determine the tone quality of a sound. Even distributions present a listener with tonal attributes such as pitch; irregular distributions create noise-like characteristics. In the figure note the complexity of the energy peaks, highly structured but irregular; also note the amount of acoustic information discarded as the same signal is reproduced at lower sampling rates. These features describe the two most elusive objectives of real-time sound synthesis: (1) to generate complex harmonic structures, (2) at high sampling rates.

Trumpet tone at decreasing Sample Rates

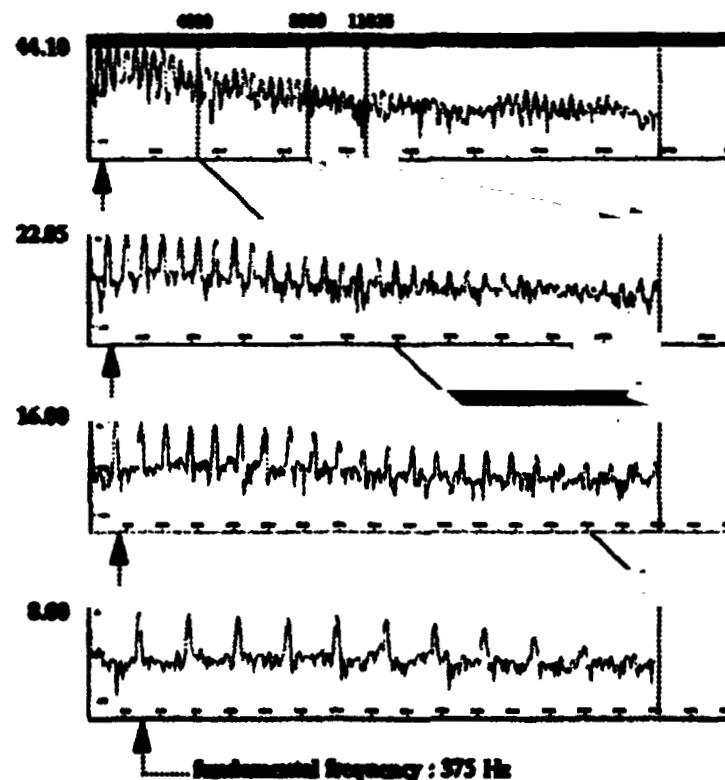
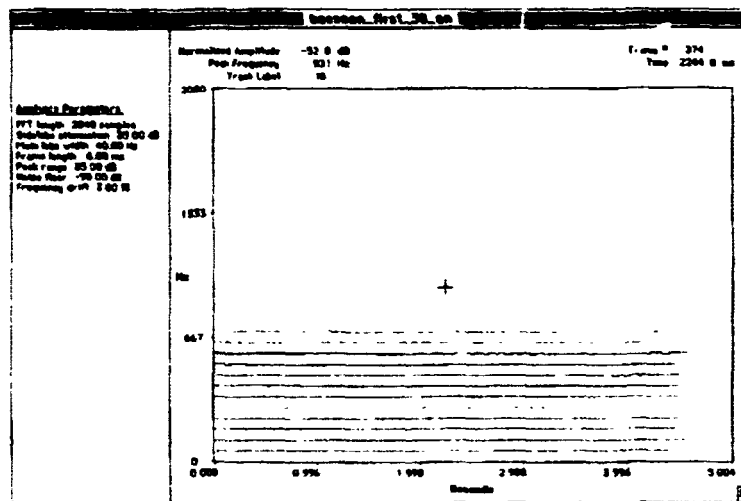


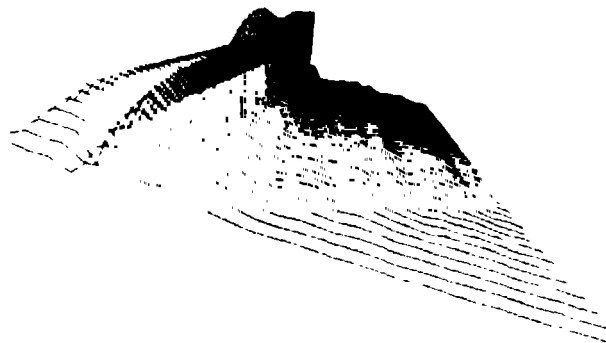
Figure 3

SOUND EVOLVING IN TIME

In natural sounds, frequency-domain structures evolve in complex ways in the time domain. This upper figure shows the energy peaks in a single bassoon tone: frequency is depicted on the vertical axis, time on the horizontal axis and amplitude by greyscale. The lower figure provides a better view of the amplitude evolution over time. The regular distribution in frequency and stability of peak locations in time indicates that the tone is quite harmonic (having a well-tuned pitch). Note even with this regularity the high degree of complex variation in local structure. The human ear is very good at comparing one such structure with another. The potential to hear distinguishing features in resonating systems at this level of acoustic structure encourages the pursuit of sonification tools for studying high-dimensional data that may have hard-to-detect regularities.



Spectral analysis of a bassoon tone. Time on the X axis, frequency on the Y axis, amplitude indicated by darkness of lines.



The same bassoon tone analysis viewed as a spectral surface, with frequency on the X axis, amplitude on the Y axis, and time receding along the Z axis.

Figure 4

SOUND SYNTHESIS ENGINES

Two taxonomies of sound synthesis methods account for the range of solutions to the problem of generating complex signals. Dodge [2] describes three broad classes: additive accumulation of simple waveforms; modulation of one waveform by another to produce sidebands; and filtering of a broadband (noisy) signal to obtain desired energy peaks. Each of these produces a steady-state waveform with controllable harmonic and noise characteristics. A waveform may be generated by continuous functions or lookup tables with a corresponding tradeoff between flexibility and computational efficiency. To obtain waveforms varying in time, additional control signals are applied to the amplitudes and frequencies of the source signals during the course of a synthesized sound. The problem of organizing the control signals in efficient and structured ways remains unsolved. Smith [3] provides a classification of synthesis strategies organized by models. The models provide varying degrees of criteria for time-domain evolution of the signal. Digitized sounds are already complex signals; it is difficult to manipulate them to produce different sounds. Spectral models organize the trajectories of energy peaks in a sound over time; analyses of natural sounds may be used to obtain guidelines for the time-based control signals that are required. Physically-based models describe coupled excitor-resonator systems with sets of ordinary differential equations. These provide efficient time and frequency descriptions; however, they are difficult to control and offer many unpredictable solutions. Smith's last category is a catch-all for systems that do not follow models based upon the reproduction of natural sounds.

Dodge:				
	Additive	Distortion	Subtractive	
Smith:				
	Processed Recording	Spectral Model	Physical Model	Abstract Algorithm

Figure 5

AN EARLY VIRTUAL ENVIRONMENT

In the early 16th century Albrecht Dürer recorded the research efforts of visual artists to harness the principles of linear perspective. Historically it is noteworthy that the artists' efforts pre-date those of geometers to understand Euclidean projection in graphical terms [4]. This etching portrays mechanisms that also operate in a VE system, particularly knowledge of the user's position and orientation with respect to other objects, and the capability to render visual information accordingly. Consider what might be the acoustic analogy to the visual systems depicted here. One analogy is *localization*, the presentation of sounds from various positions and distances measured with respect to the user. Visually there are a number of relations the artist can obtain by the use of perspective, in addition to the representation of distance. The position of the frame provides a particular discourse concerning the positions of the objects that are framed. The frame not only defines an observer's relation to the scene, it defines the relations of the objects within the scene to one another. We may ask, in sound do we have analogies to these visual *frames of reference*? One analogy is the implicit need for a model of the space shared by the listener and sound sources, a space in which the sound reverberates. Unlike light, we attend to sound simultaneously in all directions. Another analogy is the need to compare sounds with one another, to arrive at complex relations and subtle meanings such as relative degrees of importance or degrees of similarity and differences among objects which the sounds represent. Research by musicians and composers will be of great benefit to creating acoustic frames of reference in VE systems.



Albrecht Dürer: *Man Drawing Reclining Woman*. Engraving from *Measuring Instructions*, printed in Nuremberg, 1525. Kupferstichkabinett, Berlin.

Figure 6

VISUAL FRAMES OF REFERENCE

These images demonstrate the influence of neighboring patterns upon the perception of a whole. Curvilinear groups convey different messages depending upon their re-contextualization by other groups. None of the groups below have a strong representational function in isolation. Together, converging lines become a road, vertical lines become poles and curves become a mountainous horizon. Can we say absolutely that these figures do or do not convey these meanings? Together, they convey *my intention* to convey these meanings, an intention in which you participate if you also see the contents I enumerate. Again inviting analogy to sound, we wish in VE to assemble acoustic signals to convey meaningful inter-relations rather than abstract figures. Let us also understand that acoustic or visual messages do not emerge from scientific or engineering data without the presence of intentional designs to enable the assemblage of meaningful relations according to principles of perception and cognition.

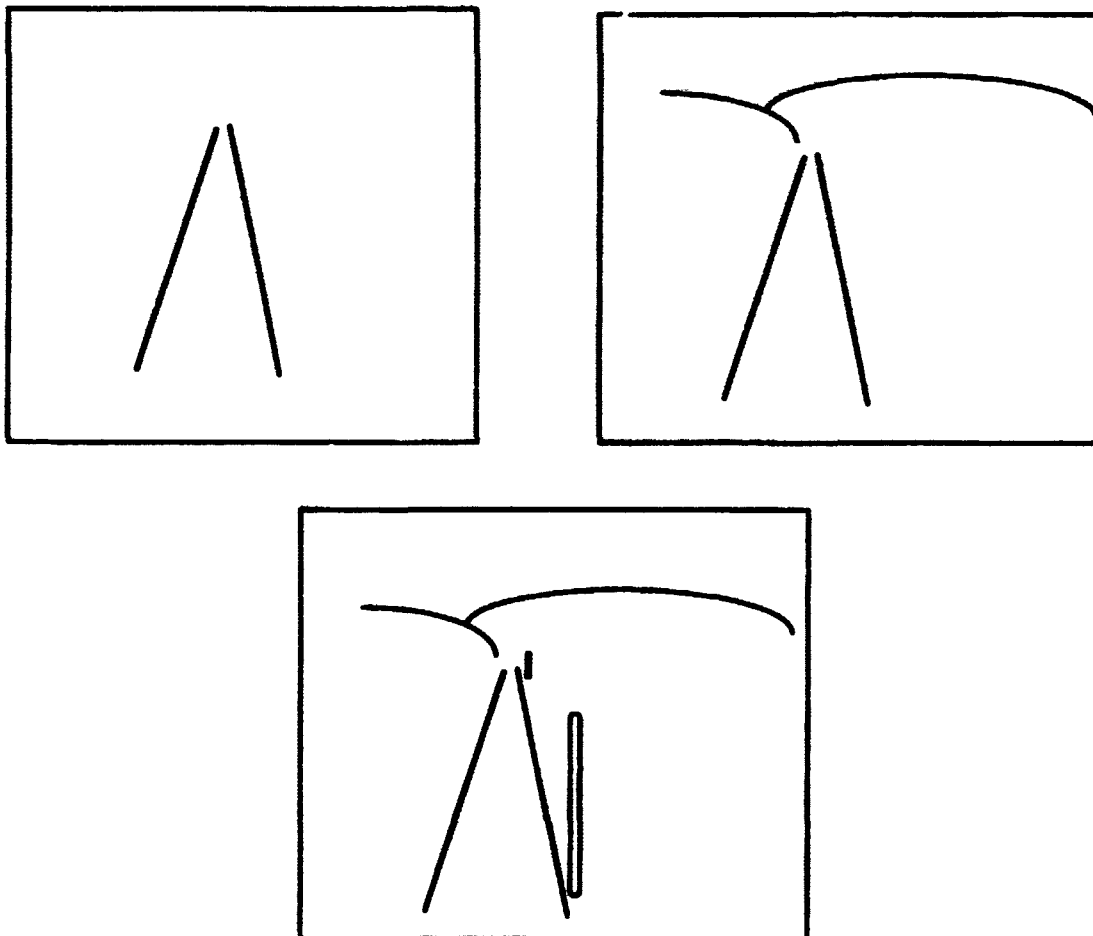


Figure 7

ACOUSTIC FRAMES OF REFERENCE

An excerpt from a string quartet from Haydn [5] provides examples of acoustic frames of reference constructed from abstract figures in sound. The musical staff orders the instruments by ascending frequency range, 'cello, viola, second violin, first violin. Vertical lines across all four parts indicate the time passing in measures. Vertical coincidence of notes indicates simultaneity. Throughout most of this example the first violin has a more active part, supported by the others making more regular sounds that change more slowly. A discourse is established in reference to a small collection of musical patterns, which may be shared among the players. Significant changes are perceived not on a note-by-note basis, but across the discourse of patterns.



Figure 8

For example, at measure 40 violin 1 ascends in an ornamented passage while the others play together in a steady pulse; at m. 42 the lower three instruments sustain single tones while violin 1 descends through the acoustic space opened up in the previous two measures. In music this *solo-accompaniment* relation is similar to visual *figure-ground* systems. A conversation begins in m. 44 as violin 2 and viola trade patterns with violin 1. The 'cello rests in mm. 45 and 46, providing a silence in the lowest frequency range. One function of silence is to emphasize a sound upon its return, such as the return in m. 47 of the lower the instruments' accompaniment role against a loftier violin 1. Another role of silence is to emphasize a sound by isolation, as violin 1 solo reaches a peak in m. 49 while the others rest. In mm. 50-53 the conversation and rests are redoubled and shared by all players, reaching a temporary conclusion and punctuation when all play and rest together. The terminal symbol on the musical staff indicates the passage will be repeated. The composer chooses repetition for structural emphasis before going on to new material.

CONSTRUCTING A MULTI-MODAL EXPERIENCE

The cinema is the dominant paradigm for audio-visual messages. This figure represents essential cinematic features: images in discrete frames that hold the screen unchanged when they are displayed, while sounds accompany the images in a continuous signal, having no notion of "frame." The dichotomy, motionless image - frameless sound, carries over into digital media, and with it come a host of complications regarding the conjunction of sound and image. Many of these complications have never been resolved in the cinema; instead, the industry adopted work-arounds that are now communications conventions. Computer-based media are capable of finding new technical solutions to image-sound incompatibilities; in so doing we may challenge existing communications conventions. Issues that arise in delivering a real-time audio-visual message stream include time-critical computing in UNIX, negotiated graceful degradation when processes overtax the CPU, separation of VE applications from an application framework, and locating outside of the application program specialized engines such as physics modules. We have already touched upon basic sound modeling; next we will discuss rendering, synchronization and display.

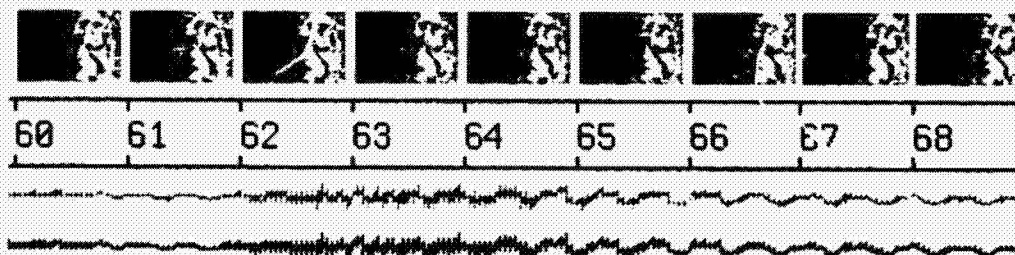


Figure 9

PARALLEL RENDERING PIPELINES

We propose an alternative to the cinematic model. In the cinema, sounds and images may be captured from anywhere and placed together on the film. In our alternative system, sounds and images come from the dynamics of a single numerical model. This rigorous restriction defines new boundaries for audio-visual communications. The knowledge that sound and image are both originating in a single source model allows experimental observations about the state of the underlying model. This sort of observation is not part of the conventional cinematic experience. Using parallel rendering pipelines we may be able to represent experimental data with cinema-like, naturalistic display strategies. Research is needed to investigate and design transfer functions for extracting control signals for image and sound rendering.

Ideal: Parallel Rendering Pipelines

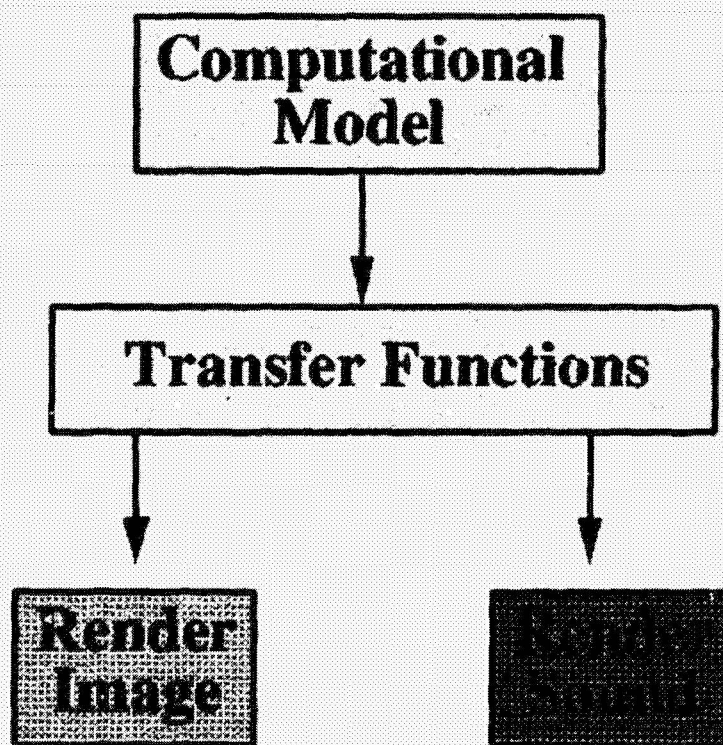
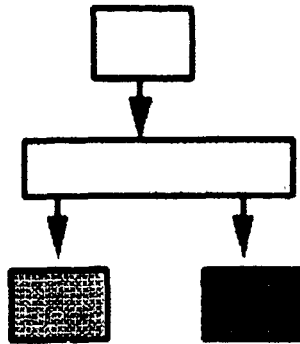


Figure 10

PARALLEL RENDERING PIPELINES: FEATURES

The capability to generate both sounds and images from a single apparatus, the computer, offers desirable features for developing robust audio-visual correlations for making experimental observations.



- **Single Hardware Platform**
- **Single OS and File System**
- **Single Programming Language**
- **(Eventually: Single Frame Rate)**
- **Timing controlled at top or bottom**

Figure 11

PARALLEL RENDERING PIPELINES: GRAPHICS

Hardware manufacturers of advanced graphics systems provide sophisticated hardware and software rendering pipelines. Many of these operations are available by simple function calls in high-level programming languages. Graphical scenes operating according to complex real-time dynamics may be rapidly prototyped.

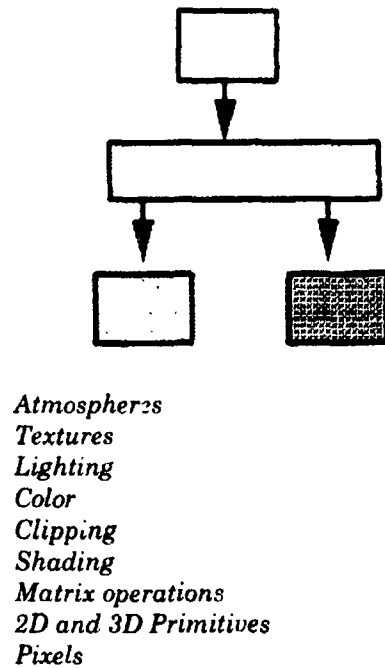


Figure 12

PARALLEL RENDERING PIPELINES: SOUNDS

If we look for hardware and software support of sound rendering on general-purpose computing platforms, we find no such architecture in existing commercial systems. High-fidelity sound rendering requires fast floating-point computation, a D/A converter and drivers, and an audio sample-buffer and scheduler protected from system interrupts. Multi-media systems on the market do not address general-purpose high-fidelity sound rendering. Multi-media systems are currently geared toward low-power desktop machines with special hardware support devices, and offer linear reproduction of sound and image sequences that were created on non-real-time platforms and are primarily non-interactive. High-level computing platforms which have the power to render sound in real-time have so far not been targeted for development of the necessary converters, drivers and libraries. Considering the capability of sound to assist in the interpretation of computations performed on powerful platforms, the lack of support for audio takes on the appearance of an oversight, or at best a lack of imagination.

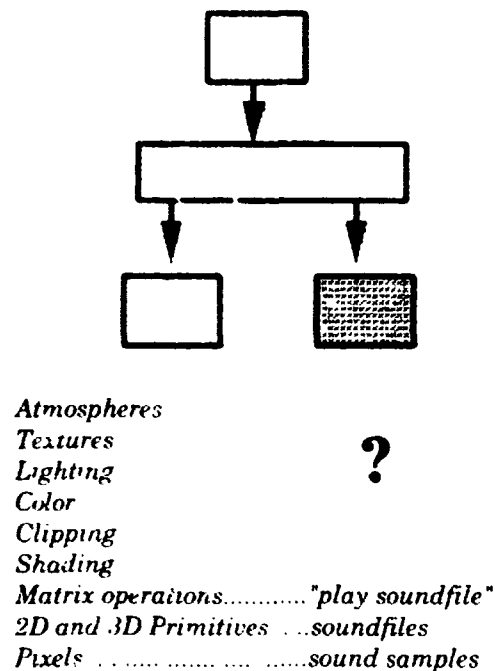


Figure 13

THE NCSA SOUND SERVER

The NCSA Audio Development Group conducts research and provides software prototypes to address the need for a real-time interactive sound rendering system to function in parallel with graphical systems. We created the NCSA Sound Server to explore the capability for sound rendering in a general-purpose computing environment [6]. The Sound Server is written in C++ and runs in UNIX, with a scheduler (HTM) optimized for high-level communications to a D/A converter architecture in real-time [7]. The Server includes libraries for sound synthesis and signal processing (VSS), and high-level "Actors" containing networks of transfer functions for translating numerical signals into intelligible acoustic patterns. Communications protocols allow our libraries to be controlled from client applications. Client and server may run on separate machines, passing messages using the serial udp protocol. An interface configuration file format allows the control of the mapping between client and server at run time. This is critical for practical purposes as it allows sound design to be located outside of the client application, increasing the likelihood of immediate interactive testing using the client as a sound controller to provide actual data conditions.

CLIENT - SERVER ARCHITECTURE NCSA SOUND SERVER

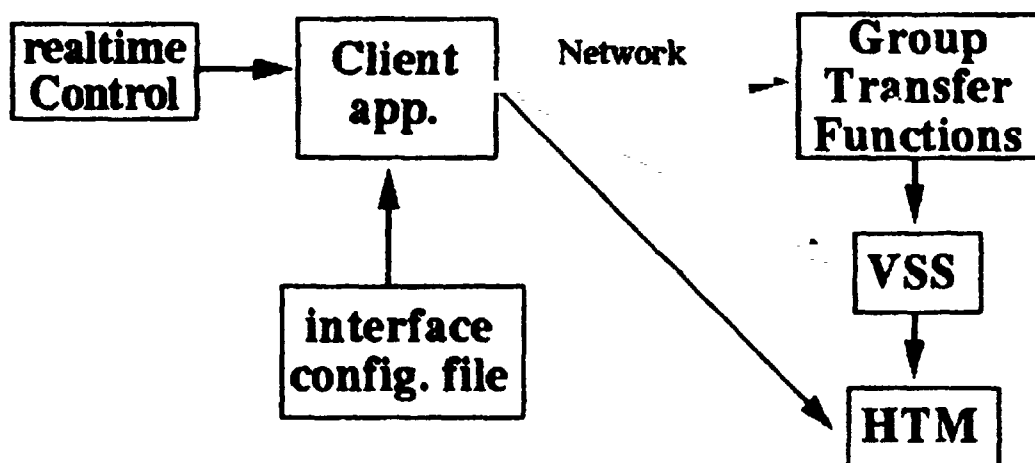


Figure 14

THE NCSA SOUND SERVER: FEATURES

Advantages typically associated with client-server architectures provide a favorable media development environment for applying sound to scientific computation.

Client-Server Advantages

- **Less code to merge - prototypes easily**
- **Audio code remains independent and stable**
- **VE client becomes synthesis interface**
- **Clients run on platforms other than SGI**
- **Sound synthesis in real-time in UNIX**

Figure 15

PARALLEL RENDERING PIPELINES IN THE CAVE

The EVL-NCSA CAVE[®] provides a testbed for applying sound rendering in parallel to graphics environments. Most CAVE applications include a computational engine that models the environment as well as the graphical rendering functions. CAVE clients link to audio libraries at compile time. The client application typically runs on a dedicated multiprocessor machine, while the sound server requires an interrupt-protected CPU and usually runs on a dedicated machine, receiving messages via serial communications.

CAVE Client – Sound Server Architecture

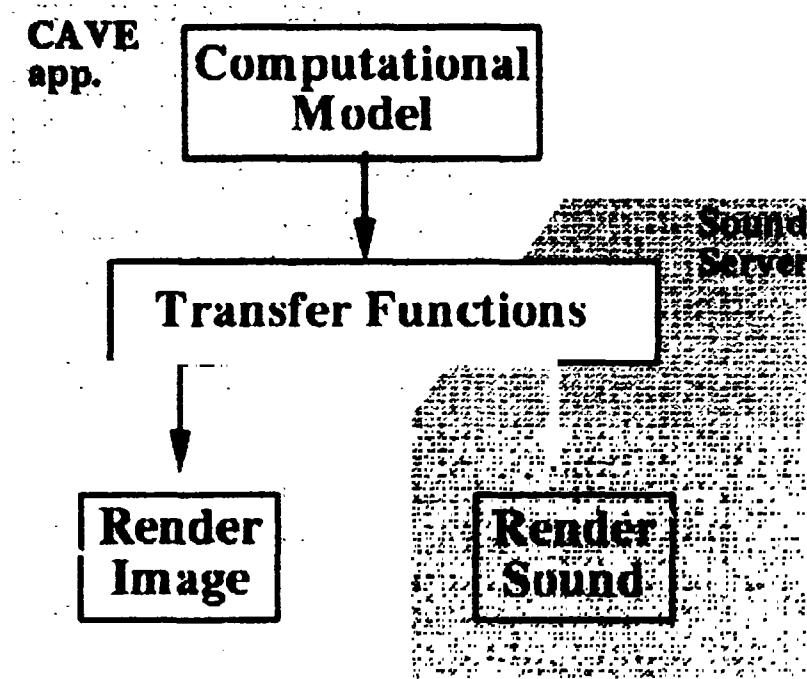


Figure 16

SOUND-IMAGE SYNCHRONIZATION: THREE HEADACHES

Three attributes of standard graphical rendering architecture contradict the needs of sound rendering systems. First, high-fidelity sound requires a sample-loop execution 48,000 per second. Graphical frame rendering loops perform at much slower rates. Second, the display rate of rendered frames is allowed to vary radically, whereas sound needs to be displayed at a constant uninterrupted sample rate. Pauses as short as two samples in duration will create noticeable discontinuities in the form of bothersome clicks in an audible signal. Third, graphical rendering pipelines have no concept of scheduling other than "next in line" and "as soon as possible." Even if visual and audible samples are rendered at the same time in their respective pipelines, there is no way to guarantee with existing hardware that the results will reach the display devices at the same time.

The Reality of Graphics Frame Rates

- **Resolution of 10-30 frames per second**
- **Vary with CPU load**
- **No concept of display time**

Figure 17

SOUND-IMAGE SYNCHRONIZATION: COMPETING DEPENDENCIES

In many virtual environments the update of the entire system is determined by the frame rate of the graphical display. This presents a problem for sound if it is to be synthesized within a graphics loop. Comparing the computation loops of graphics and sound samples we notice they operate on incompatible concepts of time. Graphical display is dependent upon upcoming events: the current frame remains on screen until the next frame is finished rendering. Display time varies accordingly. Auditory display is dependent upon passing events: the current sample buffer is displayed at a fixed sample rate and as soon as it is completed the next buffer must begin its display in order to avoid interruptions in the signal. Audio is computed in variable buffer-lengths to compensate for the fixed display rate. Human sensitivity to time discontinuities appears to be lower for visual signals than for audio signals: a noticeable variation in visual frame rate does not prohibit the interpretation of visual form and motion, whereas human perception of audio signals cannot tolerate a comparable degree of discontinuity in time without disrupting the cognitive imaging of a signal as the product of a sounding body in a real world.

Sample Computation and Buffering

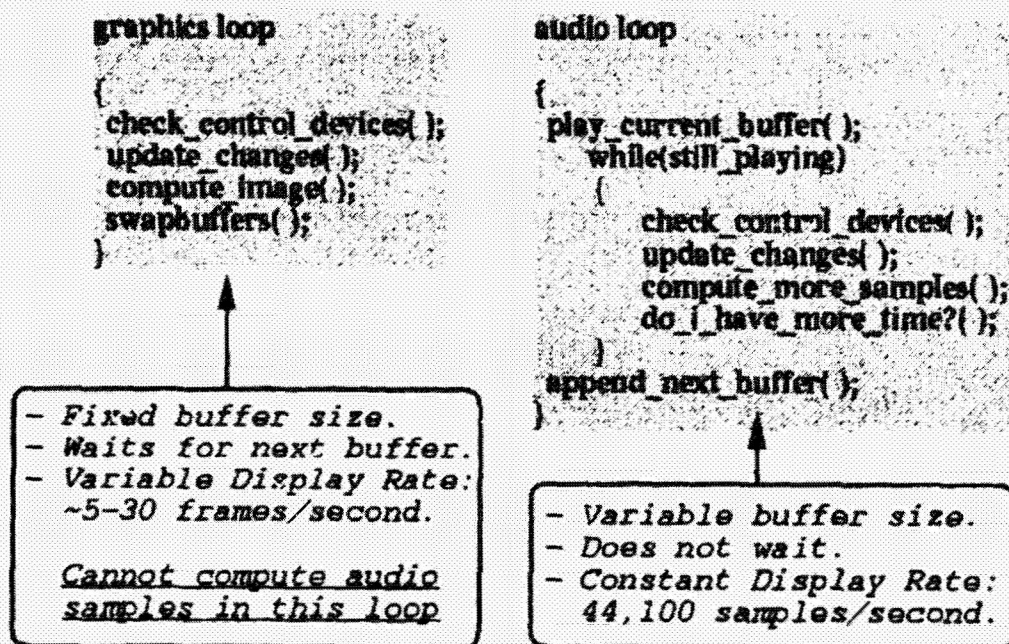


Figure 18

SOUND-IMAGE SYNCHRONIZATION: CLIENT-SERVER SOLUTIONS

A client-server paradigm permits two sample computation strategies to occur in parallel without conflicting dependencies. Sound and graphics do not share a computation loop, instead they are coordinated at two different locations: first, by sharing common data at their source; second, by high-level (but not low-level) time coordination of events. Each engine can run at its optimal rate and update under separate conditions. This requires the sound models to have sufficient intelligence to compute waveform trajectories independent of visual-based control information. Sounds update independently and receive high-level control signals from the graphical and interactive environment. These controls are generated no faster than the graphical frame rate, a good rate for phrase-level audio events as long as the integrity of the waveform evolution at 48 kHz is not interrupted. "Phrase level" events occur in sound at rates of roughly slower than 20 Hz, the rate at which a stream of changes in sound pressure level can be perceived as a steady tone.

Sample Computation and Buffering

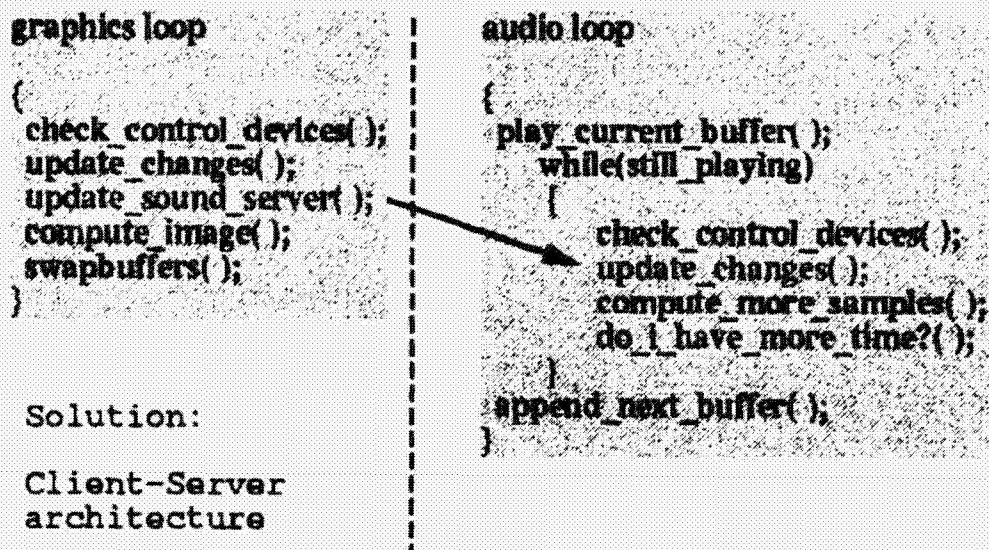


Figure 19

THE NCSA SOUND SERVER IN THE CAVE

CAVE clients run on a multiprocessor computer with special graphics hardware, while the Sound Server (VSS) runs on a separate dedicated platform with the necessary D/A conversion hardware. Downstream from the server the audio signal is multiplied in a signal matrix and signal processing is applied. In this way multiple sounds are independently localized in a 2D or 3D distribution of speakers, and distance cues (*externalization*) are applied. Positional values and moving sound sources are controlled from the CAVE application. The CPU-intensive nature of simulated localization and externalization requires dedicated hardware. This hardware is controlled from the Sound Server using the MIDI (Music Instrument Digital Interface) serial communications protocol.

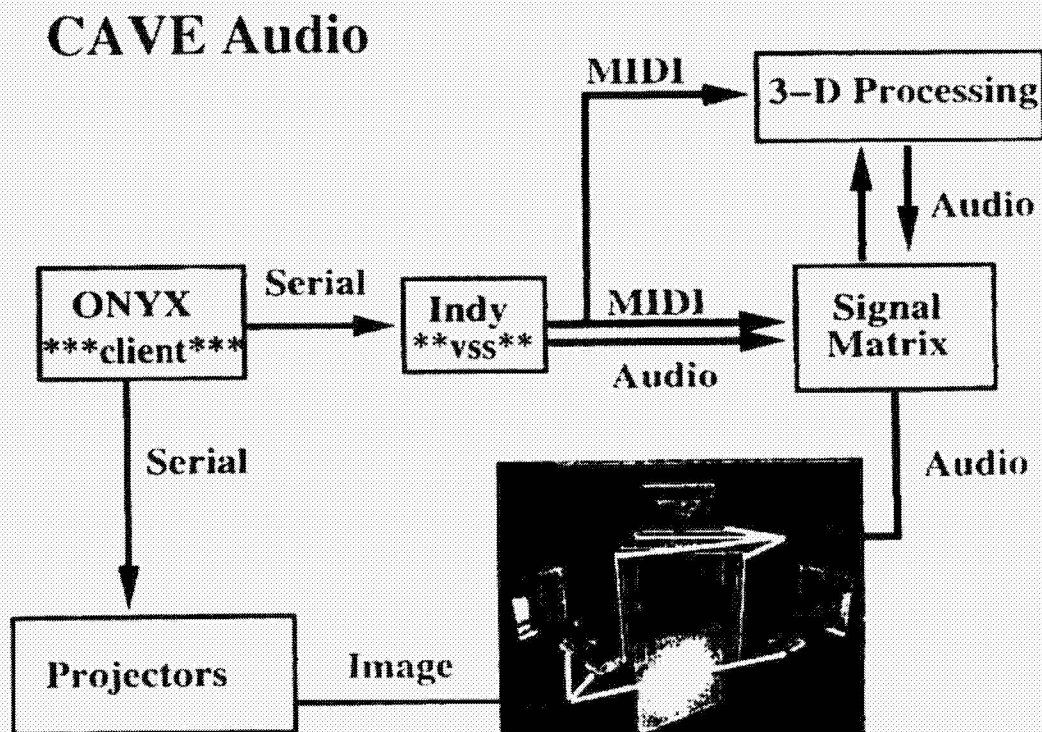


Figure 20

AN EXAMPLE APPLICATION: THE SOUND OF CHAOS

We have explored the sound of signals from the Chua's circuit, an experimental electronic circuit designed for the study of chaos [8]. In the CAVE we control a numerical simulation of the Chua's circuit with a *manifold interface* designed to allow gesture-based control of high-dimensional systems [9]. We display a graphical surface representing a control region and a cursor for navigating the surface using a gesture-based control device such as a 3D mouse or wand. In the same visual space we superimpose a phase portrait of the output signal of the three ODE's that simulate the Chua's circuit [10]. To obtain sound from the simulation the samples from one of the ODE's are sent to the Sound Server scheduler and converted directly into an audio signal. The sound changes radically during bifurcation scenarios from steady, pitched tones to regular and irregular rhythmic pulses, and then to bandpass-like noise as the state of the system moves from periodic to intermittent and chaotic regions. We cannot pass the sound samples from the CAVE client to the Sound Server in real-time at a 48 kHz rate, so we run the ODE's both in the client to obtain a visualization of the signal, and in the Sound Server (at a higher sample rate) to obtain the audible signal. The two sets of Chua's equations remain in very similar states because both are controlled in real-time by gestures from the manifold interface.



Figure 21

CONCLUSIONS

The commercial music industry offers a broad range of "plug 'n' play" hardware and software scaled to music professionals and scaled to a broad consumer market. The principles of sound synthesis utilized in these products are relevant to application in VE. However, the closed architectures used in commercial music synthesizers are prohibitive to low-level control during real-time rendering, and the algorithms and sounds themselves are not standardized from product to product. Thus a given control signal produces different results on different synthesizers. To bring sound into VE requires a new generation of open architectures designed for human-controlled performance from interfaces embedded in immersive environments.

The implementation of interactive sound synthesis in a general computing environment is a step toward "Plug 'n' Play" audio functionality in VE. Both the graphical computing and digital audio communities are just beginning to awaken to the potential needs of researchers and artists for these types of integrated tools. The NCSA Audio Group is developing high-level libraries that can be called from client applications to create well-structured audio environments. These respond to the states of a client application with special sound signals or subtle changes to the acoustic ambience in a VE display. We desire to keep our functionality in software as much we can, with obvious tradeoffs between low-level control and speed of execution. In software we have the greatest chances of developing a uniform set of protocols to be used and upgraded by the scientific computing community. Hardware manufacturers need to be encouraged to include audio hardware, device drivers and synthesis strategies as part of the standard tool set provided for scientific computing environments.

For further information regarding the NCSA Audio Development Group please visit our web page at <http://www.ncsa.uiuc.edu/VEG/audio> .

REFERENCES

1. Choi, I., "Sound Synthesis and Composition Applying Timescaling Techniques for Observing Chaotic Systems," *Proceedings of the Second International Conference on Auditory Display, ICAD '94*, G. Kramer (ed.), Sante Fe Institute, Santa Fe, 1995, pp. 79-107.
2. Dodge, C. and Jerse, T., *Computer Music: Synthesis, Composition and Performance*, Schirmer Books, NY, 1985.
3. Smith, J. O., "Viewpoints on the History of Digital Synthesis," *Proceedings of the International Computer Music Conference*, Montreal, CA, 1991, pp. 1-10.
4. Panofsky, E., "Dürer as a Mathematician," *The World of Mathematics*, J. R. Newman (ed.), Redmond, WA, Tempus, 1988, pp. 595-612; reprint of *The World of Mathematics*, Simon and Shuster, NY, 1956; original from Panofsky, E., *Albrecht Dürer*, Princeton University Press, 1943.
5. Haydn, F. J., "Finale," String Quartet Op. 74, No. 3, (Hob. III/74), 1793, *Anthology for Musical Analysis*, C. Burkhardt (ed.), Holt, Rinehart and Winston, NY, 1986, pp. 153-165.
6. Bargar, R., Choi, I., Das, S. and Goudeseune, C., "Model-Based Interactive Sound for an Immersive Virtual Environment," *Proceedings of the International Computer Music Conference*, Aarhus, Denmark, 1994, pp. 471-474.
7. Freed, A., "Tools for Rapid Prototyping of Music: Sound Synthesis Algorithms," *Proceedings of ICMC*, San Jose, Oct. 1992.
8. Chua, L. O., Wu, C. W., Huang, A. and Zhong, G. Q., "A Universal Circuit for Studying and Generating Chaos, Part I: Route to Chaos," *IEEE Transactions on Circuits and Systems - I: Fundamental Theory and Applications*, Vol. 40, No. 10, Oct. 1993.
9. Choi, I., "Interactive Exploration of a Chaotic Oscillator for Generating Musical Signals in Real-Time Concert Performance," *Special Issue: Chaos and Nonlinear Dynamics, Journal of the Franklin Institute*, June 1995.
10. Kennedy, P. M., "Three Steps to Chaos, Part II: A Chua's Circuit Primer," *IEEE Transactions on Circuits and Systems - I: Fundamental Theory and Applications*, Vol. 40, No. 10, Oct. 1993, pp. 657-674.

**NEXT
DOCUMENT**

55-71
6244

THE MANY FACETS OF AUDITORY DISPLAY

Meera M. Blattner*
University of California, Davis
Lawrence Livermore National Laboratory
Livermore, CA 94550

INTRODUCTION

It would be difficult to imagine a virtual world without sound. Sound surrounds us constantly; the ability to hear sound is one of our basic senses. Why hasn't sound become an integral part of the human-computer interface? There are many historical reasons why this is the case: the letters of the alphabet when typed into a keyboard were easily interpreted into binary form for textual displays upon a screen or printed page. Voice input has many difficulties when used as an input medium and corresponds more to the difficulties of using handwritten input, where errors occur because input is not precise. In the past, nonspeech audio has been associated with music and has not been used for conveying information, with certain exceptions such as bugle calls, fog horns, talking drums, etc., which were not universally known and limited in scope (Ref. 1). Interfaces of the future will be designed for human expression with all its subtleties and complexities (Ref. 2). One such example is the use of facial expressions used by interface agents that act as guides to assist their human users in making decisions. The range of expressiveness in audio is just starting to be appreciated by interface designers. Audio can be used to create intense emotion through music or to enhance our perception of real-world phenomena through auditory display.

In this presentation we will examine some of the ways sound can be used. We make the case that many different types of audio experiences are available to us. We should not limit our use of audio to one type of sound or even several types. A full range of audio experiences include: music, speech, real-world sounds, auditory displays, and auditory cues or messages. The technology of recreating real-world sounds through physical modeling has advanced in the past few years allowing better simulation of virtual worlds. Three-dimensional audio has further enriched our sensory experiences.

*Also with the Department of Biomathematics, M.D. Anderson Cancer Research Hospital, University of Texas Medical Center, Houston.

This work was performed with partial support of NSF Grant #IRI-9213823 and under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

SOUNDSCAPES

The computational limitations of real-time interactive computing do not meet our requirements for producing realistic images for virtual reality in a convincing manner. Regardless of these restrictions the representations can be no better than the graphics. Computer graphics is still limited in its ability to generate complex objects such as landscapes and humans. Nevertheless, useful and convincing visualizations can be made through a variety of techniques.

A similar situation is true for sound for virtual reality. It is beyond our ability to create interactive soundscapes that create a faithful reproduction of real world sound; however, by choosing one's application carefully and using sound to enhance a display rather than only mimic real-world scenes, a very effective use of sound can be made.

SOUNDSCAPES

- **We cannot create interactive soundscapes that are faithful reproductions of real world sounds.**
- **We should use sound to enhance a display rather than only mimic real-world sounds.**

Figure 1

OBJECTS

What we hear is very different from what we see. We do not always hear objects in the real world. Some objects are only heard because they are out of sight. Some objects make sounds at some times, but not others. Objects only make sounds when they set up vibrations in a medium that surrounds us, such as air or water. This means that they must create a movement to make a sound. Even movement does not always create sounds we can hear.

To associate sounds with an object that does not ordinarily make sound is artificial and there may not be natural associations with these sounds. Sound can be used very effectively to indicate the presence of objects that are not seen. Film sometimes does this through the use of music or other sound effects.

OBJECTS

- **Objects are usually seen.**
- **Some objects are only heard.**
- **Sometimes the presence of an object is detected by sound and interaction with another object**
 - ◆ **footsteps**
 - **coughing**
 - ◆ **doorbells**
 - **electrical equipment**
- **Sound may tell us how an object is constructed.**

Figure 2

MUSIC AND SPEECH

A powerful use of music is found in film scores. Music comes to bear in helping to realize the meaning of the film, in stimulating and guiding the emotional response to the visuals. Music serves as a kind of cohesive, filling in empty spaces in the action or dialogue, and the color and tone of music can give a picture more richness and vitality and pinpoint emotions and actions. It is the ability of music to influence an audience subconsciously that makes it truly valuable to the cinema. Music specific to particular cultures is used in the study of history, geography and anthropology. A scene placed in a geographical context may be enhanced by local music.

Speech is required for detailed and specific information. It is through speech (rather than through other sounds) that we communicate precise and abstract ideas. Speech may be used as input as well as output in the computer interface. Very little is known about building successful speech interfaces for two-dimensional displays, let alone three-dimensional interfaces.

TYPES OF SOUNDS

- ✓ **Music**
- ✓ **Speech**
 - Real-world sounds**
 - Auditory displays**
 - Cues and auditory messages**
 - 3D Auditory displays**

Figure 3

REAL-WORLD SOUND AND AUDITORY DISPLAYS

Real-world sounds are the natural sounds of the world around us, such as leaves rustling or birds singing, or man-made sounds such as machine noises or even a band playing in the background. What about sound in our everyday life? Real-world sounds are essential to our sense of presence in a scene that depicts our world around us. R. Murray Schafer (Ref. 3) describes "soundscapes," as historical reconstructions of the sound that surrounds people in various environments. Examples are street criers, automobiles, the crackling of candles, church bells, etc.

Auditory displays include the interpretation of data into sound, such as the association of tones with charts, graphs, algorithms or sound in scientific visualization. These auditory display techniques are used to enable the listener to picture in his or her mind real-world objects or data. An example of auditory display is the work done by Mansur, Blattner and Joy (Ref. 4), in which points on an x-y graph were translated into sonic equivalents with pitch as the x-axis and time on the y-axis (a nonlinear correction factor was used). Recently, Blattner, Greenberg, and Kamegai (Ref. 5) enhanced the turbulence of fluids with sound, where audio was tied to the various aspects of fluid flow and vortices.

TYPES OF SOUNDS

- Music**
- Speech**
- ✓ **Real-world sounds**
- ✓ **Auditory displays**
- Cues and auditory messages**

- 3D Auditory displays**

Figure 4

CUES AND AUDITORY MESSAGES

Generally audio cues will be considered auditory icons or *earcons*, to provide information to the user. This information tends to be more abstract than that received through auditory displays. Auditory signals are detected more quickly than visual signals and produce an alerting or orienting effect (Ref. 6). Nonspeech signals are used in warning systems and aircraft cockpits. Alarms and sirens fall into this category, but these have been used throughout history, long before the advent of electricity. Examples are military bugle calls, post-horns, church bells that pealed out time and the announcements of important events.

Work on *auditory icons* was done by Gaver (Ref. 7) and *earcons* by Blattner, Sumikawa, and Greenberg (Ref. 8). Gaver uses sampled real-world sounds of objects hitting, breaking, and tearing as described in real-world sounds above. However, Gaver's auditory icons are meant to convey information of a more abstract nature, such as disk errors, etc. Gaver used the term "everyday listening" to explain our familiarity with the sounds of common objects around us. Blattner, Sumikawa, and Greenberg took musical fragments, called motives, and varied their musical parameters to obtain a variety of related sounds. We describe the construction of earcons below.

TYPES OF SOUNDS

✓ Cues and Auditory Messages

- Auditory messages or signals were used by people before the discovery of electricity.
- Bells, bugles, trumpets and drums served information to the countryside or announced the arrival of an important person.

Figure 5

VIRTUAL REALITY, TELEPRESENCE AND TELECONFERENCING

The types of sounds described above, speech, music, audio cues, and real-world sounds, can all be located in a three-dimensional audio environment. Sound localization by NASA has shown the effectiveness of separating voices in space to improve their clarity (Ref. 6). Cohen and Wenzel (Ref. 9) are studying the three-dimensional acoustic properties of teleconferencing systems to filter out extraneous sounds by the use of "audio windows." The general idea is to permit multiple simultaneous audio sources, such as in a teleconference, to coexist in a user-controlled display to easily move through the display and separate the channels while retaining the clarity and purity of the sounds.

TYPES OF SOUNDS

- ✓ 3D Audio
- Virtual Reality
- Telepresence
- Teleconferencing

Figure 6

THREE-DIMENSIONAL SOUND

Three-dimensional (localized) sound truly immerses the listener in his or her auditory environment. The basis of the work in three-dimensional acoustic displays is psychoacoustics. The virtual acoustic environment is part of the NASA Ames View System (Ref. 6). The technology of simulating three-dimensional sound depends on reconstructing the sounds as they enter the ears. The acoustic signals are affected by the pinnae (outer ear) and the distance and direction of the ears. Microphones were placed in the ears of humans or mannequins to measure this effect, called the head-related transfer function. A real-time system, the Convolvotron, is used to filter incoming sounds using a head-related transfer function (Ref. 6).

3-D AUDITORY DISPLAY

Synthesis Technique

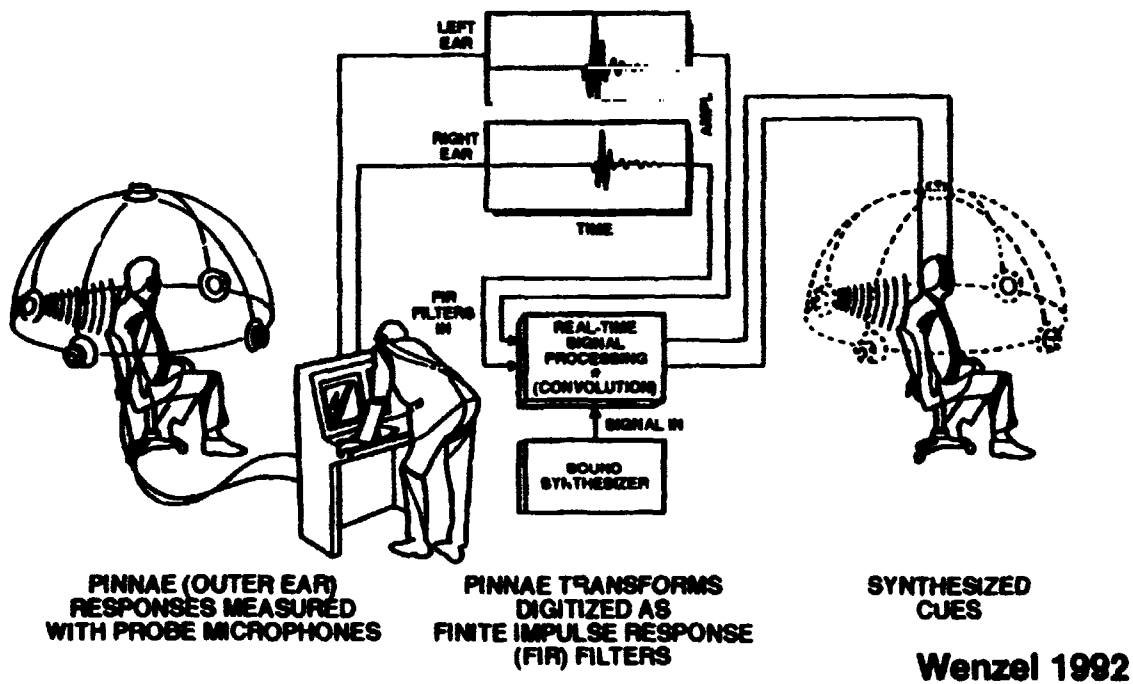


Figure 7

PARAMETERS OF SOUND

Audio has dimensions or parameters. In nonspeech audio these parameters are manipulated to provide the symbols or syntax of messages. The dimensions of sound are (Ref. 9):

- harmonic content
 - pitch and register (tone, melody, harmony)
 - wave shape (sawtooth, square, ...)
 - timbre, filters, vibrato, and equalization
- dynamics
 - intensity/volume/loudness
 - envelope (attack, decay, sustain, release)
- timing
 - duration, tempo, repetition rate, duty cycle, rhythm, syncopation
- spatial location
 - direction (azimuth, elevation)
 - distance/range
- ambience: presence, resonance reverberance, spaciousness
- representationalism: literal, abstract, mixed.

THE PARAMETERS OF SOUND

■ Graphical parameters are (Ref. 10):

- size
- saturation
- texture
- orientation
- shape
- color

■ Sound parameters are (Ref. 9):

- harmonic content
- dynamics
- timing
- spatial location
- ambience
- representationalism

Figure 8

SAMPLED VERSUS DIGITIZED SOUND

Sampled sounds are digital recordings of sounds which we can hear. These sounds have the advantage of immediate recognizability and ease of implementation into computer interfaces. Synthesized sounds are those sounds which are created algorithmically on a computer. They can be made to sound similar to real-world sounds through sound analysis (such as Fourier analysis) and trial-and-error methods. Since synthesized sounds are created algorithmically, it is easy to modify such a sound in real time by altering attributes like amplitude (volume), frequency (pitch), or the basic waveform function (timbre). Furthermore, it is easy to add modulation of amplitude or frequency in real time to create the effects of vibrato or tremolo without changing the basic sound. It is for these reasons that sound synthesis is so popular in music creation today. Synthesized sounds offer a high degree of flexibility with a reasonable amount of ease. A drawback of synthesized sound is that each algorithm used typically mimics some sounds very well and others not as well.

Since sampled sounds are digital recordings, they can reproduce with extremely high accuracy any sound which can be heard. However, the amount of work required to attain equal flexibility in modification, compared with synthesized sounds, is very high. Typically, sampled sounds are modified only in amplitude (volume) and frequency (pitch).

SOUNDS ON COMPUTERS

- **Sampled sounds are digital recordings.**
- **Synthesized sounds are sounds which are created algorithmically.**
- **Synthesized sounds may be modified in real time by altering attributes like amplitude (volume), frequency (pitch), or the basic waveform function (timbre).**
- **Typically, sampled sounds are modified only in amplitude (volume) and frequency (pitch).**

Figure 9

THE STRUCTURE OF AUDIO MESSAGES

Earcons are short, distinctive audio patterns to which arbitrary definitions are assigned. They can be modified in various ways to assume different but related meanings. The building blocks for earcons are short sequences of tones called motives. From motives we can build larger units by varying musical parameters. The advantage of these constructions is that the musical parameters of rhythm, pitch, timbre, dynamics (loudness) and register can be easily manipulated. The motives can be combined, transformed, or inherited to form more complex structures. The motives and their compounded forms are called earcons. However, earcons can be any auditory message, such as real-world sounds, single notes, or sampled sounds of musical instruments.

EARCONS - SOUND MESSAGES

- **Motives: The basic melodic and rhythmic units**
- **A motive is either a single pitch or a sequence of two to four pitches**
- **The family motive is the specific durational sequence (rhythm) associated with the motive**
- **A motive has variable parameters of**
 - **timbre (tone color)**
 - **dynamics (loudness) and**
 - **register (high/low pitches)**

Figure 10

EARCON CONSTRUCTION

A motive may be an earcon or it may be part of a compounded earcon. Let A and B be earcons that represent different messages. A and B can be combined by juxtaposing A and B to form a third earcon AB. Earcon A may be transformed into earcon B by a modification in the construction of A. For example, if A is an earcon, a new earcon can be formed by changing some parameter in A to obtain B, such as the pitch in one of its notes. A family of earcons may have an inherited structure, where a family motive, A, is an unpitched rhythm of not more than five notes and is used to define a family of messages. The family motive is elaborated by the addition of a musical parameter, such as pitch ($A + p = B$) and then preceded by the family motive to form a new earcon, AB. Hence, the earcon has two distinct components, an unpitched motive followed by a pitched motive with the same rhythm. A third earcon, ABC, can be constructed by adding a third motive, C, with both the pitch and rhythm of the second motive, but now has an easily recognizable timbre ($A + p + t = B + t = C$).

CONSTRUCTING EARCONS

■ Combining

- The process of combining to create an earcon means linking different motives together in a chain-like sequence.

■ Transforming

- The process of transformation cosmetically alters a motive by changing its timbre, register, and/or tempo.

■ Inheriting

- The process of inheriting is one in which a single earcon is heard in an increasing complex chain.

Figure 1'

EARCON CONSTRUCTION

Inheriting earcons

Error messages for novice users

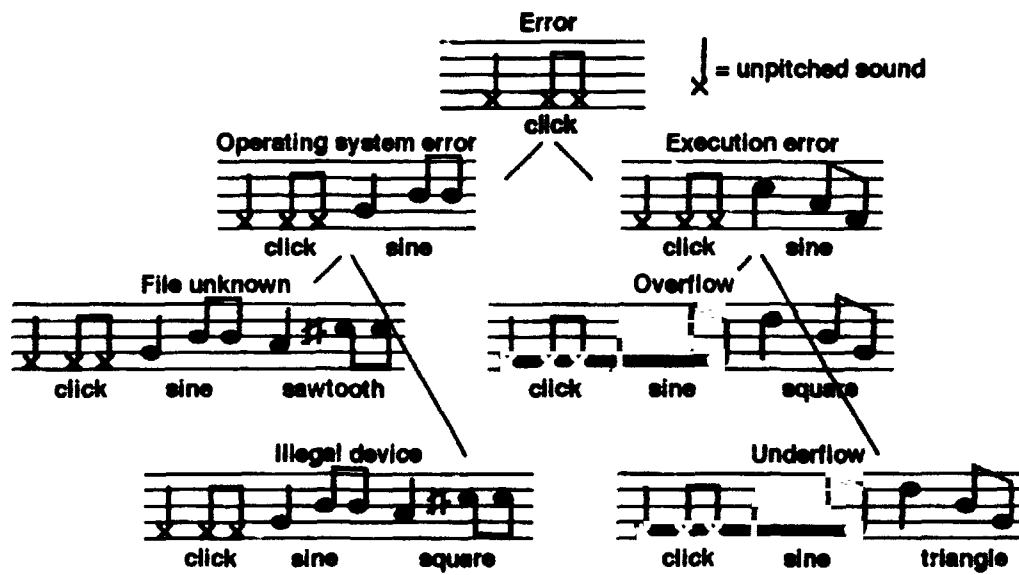


Figure 12

MULTIPLE EARCONS

To display more than one earcon their temporal locations with respect to each other have to be identified. Two primary methods are used: overlaying one earcon on top of another and the sequencing of earcons (Ref. 11). Some sort of merging or melding in to new sound could be considered; for example, the pitch of two notes can be combined into a third pitch. Programs typically play audio without regard to the overall auditory system state. As a result, voices may be played simultaneously or occur with several nonspeech messages making the auditory display incoherent. An audio server is being constructed that blends the sounds of voice, earcons, music, and real-world sounds in a way that will make each auditory output intelligible (Ref. 12).

MULTIPLE EARCONS

How to combine earcons with each other?

Two primary methods are considered here:

- **Overlaying one earcon on another.**
- **The sequencing of earcons.**
- **Some sort of merging into a new sound.**

Figure 13

WILL AUDIO MESSAGES BE USED?

Will sounds as abstract as earcons be accepted by the majority of users? The advantages are very clear: they are easily constructed on almost any type of workstation or personal computer. The sounds do not have to correspond to the objects they represent, so objects that either make no sound or an unpleasant sound still can be represented by earcons without further explanation. Auditory icons that make real-world sounds usually can be recognized quickly; however, most messages do not have appropriate iconic images.

Brewster, Wright, and Edwards (Ref. 13) found earcons to be an effective form of auditory communication. They recommended six basic changes in earcon form to make them more easily recognizable by users. These changes were: 1) use synthesized musical timbres, 2) pitch changes are most effective when used with rhythm changes, 3) changes in register should be several octaves, 4) rhythm changes must be as different as possible, 5) intensity levels must be kept close, and 6) successive earcons should have a gap between them. Earcons are necessarily short because they must be learned and understood quickly. Earcons were designed to take advantage of chunking mechanisms and hierarchical structures that favor retention in human memory. Furthermore, they are recognition rather than recall. If earcons are to be used by the majority of computer users, they must be learned and understood as quickly as possible taking advantage of all techniques that may help the user recognize them.

EFFECTIVENESS

- Will sounds that convey information in a form as abstract as earcons be accepted by the majority of users?
- The advantages:
 - ✓ Easily constructed
 - ✓ Do not have to correspond to the objects they audify
- Brewster, Wright and Edwards found earcons to be an effective form of auditory communication.

Figure 14

MAPS

To test our theories, we chose to combine auditory display techniques with two-dimensional maps. Maps are primarily used for orientation, navigation within, and analysis of, geographic terrain. However, a broad range of additional information may be of interest in some cases: average annual rainfall, soil composition, location of mineral deposits and other natural resources, location of rail lines, location of historical sights, various economic factors, elevation, etc.

MAPS

■ Maps are used for orientation and navigation.

■ Other information of interest:

- average annual rainfall
- soil composition
- location of mineral deposits
- location of rail lines
- location of historical sights, elevation
- ownership
- utilities

Figure 15

MAPS

Because they need not be static, computerized maps can take advantage of many more methods for displaying information than can traditional paper maps. Enlarged windows can appear at a point of interest; numerical data can pop up on demand and disappear when no longer needed. Animation and pseudocolor can be used to track or call attention to specific information. Nevertheless, because the addition of visual data requires that space be allocated for it, a saturation point will eventually be reached beyond which interference with text and graphics already on display cancels without any possible benefit. In such cases (and others), it may be advantageous to present some of the data in a sonic representation. Auditory maps were used by Kramer to enhance Magellan's view of Venus. Auditory output to convey information such as the emissivity (i.e., radiation) and gravity of the area being viewed. The auditory output did not disrupt the view of the underlying landscape.

COMPUTERIZED MAPS

- **Computerized maps can take advantage of many more methods for displaying information.**
- **Enlarged windows can appear at a point of interest.**
- **Numerical data can pop up on demand and disappear when no longer needed.**
- **A saturation point will eventually be reached with text and graphics.**

Figure 16

MAP IMPLEMENTATION

The floor plans are visible, as are geographical data such as roads, parking lots, etc. Information associated with each building includes sewer lines, water lines, power lines, number of computers and people housed within, the department or administrative unit in charge, construction type, level of security clearance required to work there, job titles of those in the building, etc. As the cursor is dragged over the image, relevant information is presented to the user. The mouse may remain still and information requested by the user clicking on an appropriate button. Sonic information must be presented in a "short form" when the mouse is in motion but can be presented in a "long form" when it is stationary. The short form cannot encode sufficient information to distinguish between items within a family, whereas the long form can easily do so. The functionality required involved retrieval not only of the location of particular items, but also of area information. We also needed a way to provide summary data to users. It would be too slow and inefficient to scan an entire scene with a mouse! So we had to develop methods for scanning areas and presenting multiple data. Summary data is used to indicate that there were many items of a certain type in an area. We chose a simple method to handle summary data - by a linear mapping of earcon volume (loudness) to the magnitude of the numeric value.

AUDITORY MAPS

- **An experimental system was implemented on a Silicon Graphics INDIGO workstation, "visual" cartographic data were changed into sonic representations.**
- **The data used was a map of Lawrence Livermore National Laboratory.**
- **This particular map was selected because it was in the form of a machine-readable data base of buildings.**

Figure 17

EARCON IMPLEMENTATION OF MAPS

We used timbres of various musical instruments to create earcons to help differentiate the sounds, as suggested by Brewster, Wright, and Edwards (Ref. 13). The sampled waveforms are altered by varying the frequency multiplier and the amplitude factor, to create different pitches and volumes. The frequencies range from 100 Hz to 2,000 Hz. None of the earcons varied in loudness within themselves. Instead, we interpreted summary data using dynamics; that is, sets of data with more like items are louder than sets with smaller amounts of like data.

AUDITORY MAPS

- **Timbres of musical instruments for earcons help to differentiate the sounds.**
- **TRANSFORMATIONS:**
 - **Knocking earcons are administrative access.**
 - **The computer earcon are earcons transformed over the x-axis, y-axis, and both axes with a timbre change.**

Figure 18

EARCON IMPLEMENTATION OF MAPS

Simple sounds as well as complex ones built by combination, transformation, and inheritance from motives and other earcons were used in auditory maps. There is a simple earcon of a tom-tom drum pounding, which represents a building restriction. The pounding is transformed in pitch and frequency to indicate different access restriction levels of buildings. Higher restrictions are represented by faster, higher-pitched knocking. A simple three note earcon in one pitch using a saxophone timbre indicates all properties of the first earcon except that they have different pitch changes. It is important to note that each family of earcons shares the same timbre. Since timbre is one of the most easily recognized attributes of sound, one can immediately identify the family of an earcon just by recognizing what instrument is used in playing that earcon. A combination of these different earcons can be used to build new, more complex, earcons. For instance, a combination of three earcons can indicate a physics buildings with a clearance level of confidential which houses Sun computers.

AUDITORY MAPS

■ INHERITANCE:

A pitchless earcon indicates an administrative building. The second level motives inherit all the properties of the first earcon except that they have pitch and timbre. The pitch is the same, but the timbre varies.

Figure 19

TEMPORAL QUALITIES

Sequential combinations are sounds which are heard one after another. Concurrent sounds are those which are either played simultaneously or which partially overlap in time. Combined sounds are those whose attributes are combined into a single new sonic item. An approach that is used to combine data is to consider every mapping from a data item to a dimensional coordinate system, where the coordinates are sonic attributes. The user may choose to listen concurrently or sequentially. However, if more than four earcons need to be sounded in concurrent mode, then the first four will play concurrently after which as each earcon ends another will begin. We have not implemented combined data at this time.

TEMPORAL COMBINATIONS

- **The user has a choice of listening concurrently or sequentially.**
- **If more than four earcons are displayed in the concurrent mode then four will play, and as each one ends another begins.**

Figure 20

SPEED SELECTION

Audio information can be displayed in two modes: moving and stationary. When the user moves the cursor over a region there is not sufficient time to play its earcon. A series of short, truncated sounds inform the user that there are items of interest in that location. Stationary mode is indicated when the mouse is clicked and the long form of the data under the cursor can be displayed. Selection while moving can be turned off if the user wishes.

SPEED SELECTION

■ MOVING

When the cursor is moving there is not sufficient time to play its earcon. Short sounds inform the user that there are items of interest in that location.

■ STATIONARY

Stationary mode is indicated when the cursor is clicked and then the long form can be displayed. Moving has two modes: on and off.

Figure 21

A CENTRALIZED AUDIO PRESENTATION SYSTEM

User interfaces which support concurrent program executions have little, if any, audio management. Typically, a number of audio channels exist and programs request the number of audio channels required. The operating system either grants or denies the request. Therefore, in environments where multiple programs output sound, each individual program has no overall context of the auditory systems state with the possible exception of how many audio channels have been allocated. Programs typically play audio without regard for the overall auditory environment which can cause sound masking and perceptual unintelligibility.

A CENTRALIZED AUDIO PRESENTATION SYSTEM

- **Motivation**
 - Maintain the intended informational encoding
 - Perceptual issues can then be addressed.
 - Simultaneous speech and/or non-speech audio presentation
 - Maximize clarity of each request
- **Multiple representations in audio requests**
 - Abstract earcons
 - Representational earcons
 - Voice
 - Sampled sound (no semantic content)
 - Other representations

Figure 22

THE PRESENTATION MANAGER

The presentation manager receives descriptive messages which contain information about system activities and program states, as specified by the user or application programmer. The sonic output of the set of running programs and the overall auditory system state is controlled by the presentation manager. It chooses how the information is to be presented in sound, within the constraints of the descriptive message. The presentation manager must choose the form with consideration for other current output.

MOTIVATION

- Maintain the intended informational encoding
- Perceptual issues can then be addressed
 - Simultaneous speech and/or non-speech audio presentation
 - Maximize clarity of each request

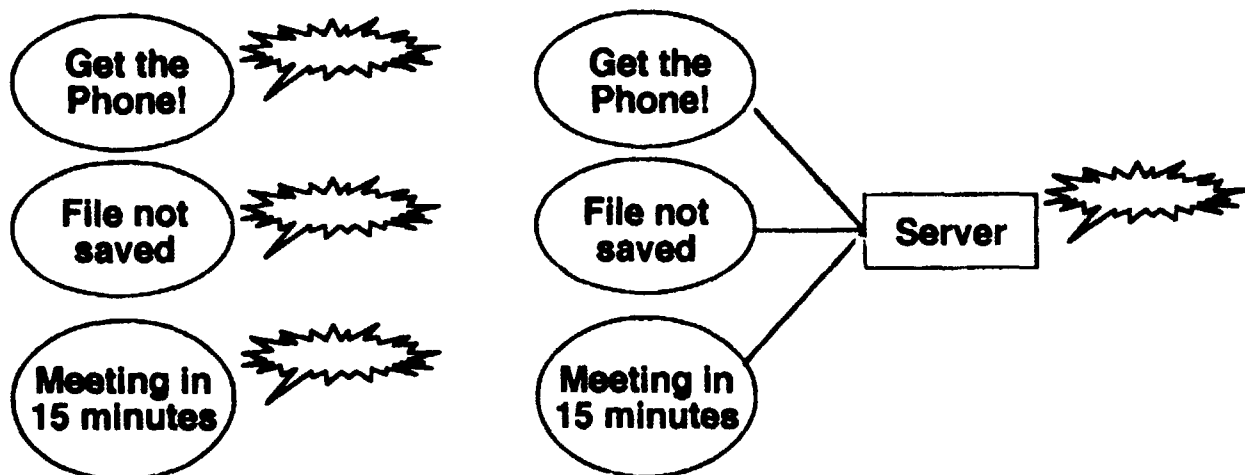


Figure 23

PRESENTATION MANAGER DESIGN

The audio presentation manager is composed of three distinct parts: the descriptive message server, the medial selector, and the scheduler. A message passing paradigm serves as the underlying model for communication between application and the various parts of the presentation manager.

Whenever an application is to represent some information in sound, it sends a message to the presentation manager. This message includes a high level description of the information to be displayed. The presentation manager then decides how the message is to be displayed.

SYNTHESIZER MODULES

- **Currently available:**
 - Earcon synthesizer
 - Voice synthesizer
 - Sine wave synthesizer
 - “Sample” player
- **In production:**
 - Algorithmic music synthesizer

Figure 24

REQUIREMENTS FOR A SYNTHESIZER MODULE

In order to give the presentation manager as much flexibility as possible, applications need not send raw auditory information. Rather, many common forms of data can be sent to the presentation manager along with more general information about that data. Depending upon the information received and the current auditory state of the system, the most appropriate auditory representation for the data will be used in its presentation.

REQUIREMENTS FOR A SYNTHESIZER MODULE

- A set of variables which constitutes its “state”
- An initialization routine for the state variables
- Must be able to compute the next n samples from the current state, and this computation must occur within $n/sample_rate$ seconds
- The ability to algorithmically encode its impact on the other forms
 - Do sounds produced by this synthesizer interfere with the perception of other sounds?
 - Can this form be sounded simultaneously with itself?

Figure 25

SERVER DECISIONS

Audio request contains many parameters to guide the decisions of the server

- Priority, latency; Is it interruptible?

- Semantics upon interruption

 - Re-play the whole request

 - Continue from where the request was interrupted

 - Remove the request completely

- Desirability of each audio form

- Information specific to the different synthesizer modules

How to determine which requests to play now:

- Function of priority, latency, current audio system state, and forms for presentation

- Must be above a minimum threshold or else that request is postponed

How to determine which form for each sounded request:

- What forms are already playing? Penalties associated with multiple forms (i.e., two or more voices)

- Which forms are preferable to the application?

- Which forms are preferable to the user in general?

SERVER DECISIONS

- **Audio request contains many parameters to guide the decisions of the server.**
- **How to determine which requests to play now.**
- **How to determine which form for each sounded request.**

Figure 26

AUDIO EXAMPLE

We implemented a simple navigation system that
Saxophone--crossed intersection
Rising trumpet--north
Drums--south

Speed indicated by beeps

Approaching intersection
Hear choices
Suggests which way to go
Request when crossing intersection
Tone indicates distance from the finish -- low to high

EXAMPLE APPLICATION

- With a strong voice user bias
- With a strong abstract earcon bias
- With no strong biases

Figure 27

SUMMARY

- **Audio is richer in three-dimensions.**
- **Sound sources are clearer when separated in space.**
- **The sense of immersion in an artificial world is greater when sound surrounds the listener.**
- **Sound can be used to replace touch.**
- **Sound can impart abstract information.**

Figure 28

REFERENCES

1. Blattner, M. M., and Greenberg, R. M., *Communicating and Learning Through Nonspeech Audio, Multimedia Interface Design in Education*, Springer Verlag, NATO ASI, Series F, 1992.
2. Blattner, M. M., "In Our Image: Interface Design in the 1990s," *IEEE Multimedia*, Vol. 1, No. 1, 1994, pp. 25-36.
3. Schafer, R. M., *The Tuning of the World*, Alfred Knopf, New York, 1977.
4. Mansur, D.L., Blattner, M.M. and Joy, K.I., "Sound Graphs: A Numerical Data Analysis Method for the Blind," *Journal of Medical Systems*, Vol. 9, No. 3, 1985, pp. 163-174.
5. Blattner, M. M., Greenberg, R. M. and Kamegai, M., "Listening to Turbulence: An Example of Scientific Audiolization," *Multimedia Interface Design*, Blattner and Dannenberg (eds.), ACM Press/Addison-Wesley, Reading, MA, 1992, pp. 87-104.
6. Wenzel, E. M., "Three-Dimensional Virtual Acoustic Displays," *Multimedia Interface Design*, Blattner and Dannenberg (eds.), ACM Press/Addison-Wesley, Reading, MA, 1992, pp. 257-288.
7. Gaver, W. W., "Auditory Icons: Using Sound in Computer Interfaces," *Human-Computer Interaction*, Vol. 2, 1986, pp. 167-177.
8. Blattner, M. M., Sumikawa, D. A. and Greenberg, R.M., "Earcons and Icons: Their Structure and Common Design Principles," *Human-Computer Interaction*, Vol. 4, No. 1, 1989, pp. 11-44.
9. Cohen, M. and Wenzel, E. A., "The Design of Multidimensional Sound Interfaces," *Virtual Environments and Advanced Interface Design*, Barfield and Furness (eds.), Oxford University Press (to appear).
10. Bertin, J., *Semiology of Graphics*, University of Wisconsin Press, Madison, WI, 1983.
11. Blattner, M. M., Glinert, E. P. and Papp, A. L., II, "Sonic Enhancement of Two-Dimensional Graphic Displays," *Auditory Display: Sonification, Audification and Auditory Interfaces*, G. Kramer (ed.), Addison-Wesley, Santa Fe Institute, Reading, MA, 1994, pp. 447-460.
12. Blattner, M. M. and Papp, A. L., III, "A Centralized Audio Presentation Manager," *Proceedings of the International Conference on Auditory Displays*, Santa Fe Institute, 1995, pp. 231-240.
13. Brewster, S. A., Wright, P.C. and Edwards, A. D. N., "An Evaluation of Earcons for Use in Auditory Human-Computer Interfaces," *Proceedings of INTERCHI '93*, Amsterdam, 1993.

**NEXT
DOCUMENT**

54
6045

HAPTIC INTERFACES: HARDWARE, SOFTWARE AND HUMAN PERFORMANCE¹

Mandayam A. Srinivasan
Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139

ABSTRACT

Virtual environments are computer-generated synthetic environments with which a human user can interact to perform a wide variety of perceptual and motor tasks. At present, most of the virtual environment systems engage only the visual and auditory senses, and not the haptic sensorimotor system that conveys the sense of touch and feel of objects in the environment. Computer keyboards, mice and trackballs constitute relatively simple haptic interfaces. Gloves and exoskeletons that track hand postures have more interaction capabilities and are available in the market. Although desktop and wearable force-reflecting devices have been built and implemented in research laboratories, the current capabilities of such devices are quite limited. To realize the full promise of virtual environments and teleoperation of remote systems, further developments of haptic interfaces are critical.

In this paper, the status and research needs in human haptics, technology development and interactions between the two are described. In particular, the excellent performance characteristics of Phantom, a haptic interface recently developed at MIT, are highlighted. Realistic sensations of single point of contact interactions with objects of variable geometry (e.g., smooth, textured, polyhedral) and material properties (e.g., friction, impedance) in the context of a variety of tasks (e.g., needle biopsy, switch panels) achieved through this device are described and the associated issues in haptic rendering are discussed.

¹Part of the work reported here was supported by Naval Training Systems Center Contract No. N61339-94-C-0087.

APPLICATIONS OF SYNTHETIC ENVIRONMENTS (SE)

Synthetic environments (SE), which include both virtual environments (VE) and teleoperation, have generated considerable excitement, owing to the wide variety of applications in which they can play a significant role (listed below; Ref. 1). At present, most of the VE systems engage only the visual and auditory senses, and not the haptic sensorimotor system that conveys the sense of touch and feel of objects in the environment. Manual interactions with SE are important in sensorimotor tasks such as training of surgeons with VE or conveying the feel of an object to the participants of a teleconference. They may also play a significant role in cognitive tasks such as memorization and analysis of multidimensional databases or teaching the implications of the violations of physical laws.

- **Design, Manufacturing, and Marketing**
- **Medicine and Health Care**
- **Teleoperation for Hazardous Operations**
- **Training**
- **Education**
- **Entertainment**
- **Information Visualization**
- **Telecommunications**

Manual interaction with SE is important in particular tasks within each application area.

Figure 1

HAPTICS: MANUAL INTERACTIONS WITH THE ENVIRONMENT - FOR EXPLORATION OR MANIPULATION

The term Haptics refers to manual interactions with real or virtual environments. It includes both exploration of objects to obtain information about the environment and manipulation of objects to alter the environment. In contrast to the purely sensory nature of vision and audition, the human haptic system involves tight integration of both sensory and motor components. Further, the sensory information can be divided into two classes: (1) tactile information, referring to the sense of contact with the object, mediated by the responses of low-threshold mechanoreceptors innervating the skin (e.g., the finger pad) within and around the contact region; and (2) kinesthetic information, referring to the sense of position and motion of limbs along with the associated forces, conveyed by the sensory receptors in the skin around the joints, joint capsules, tendons, and muscles, together with neural signals derived from motor commands.

Haptics: Manual interactions with the environment - for exploration or manipulation

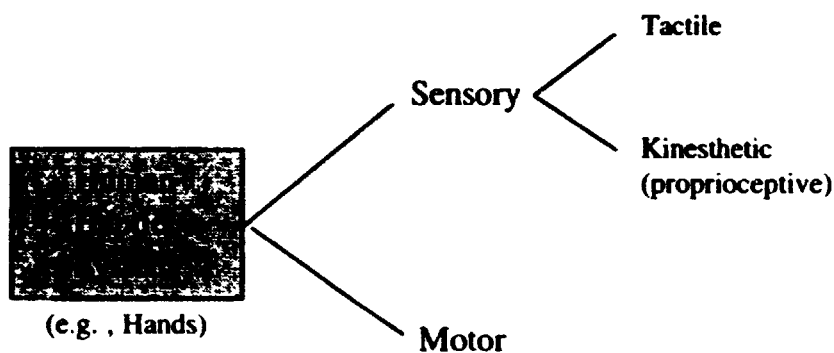


Figure 2

HAPTIC INTERFACES

Haptic interfaces are devices that enable manual interaction with virtual environments or teleoperated remote systems. They are employed for tasks that are usually performed using hands in the real world, such as manual exploration and manipulation of objects. In general, they receive motor action commands from the human user and display appropriate tactual images to the human. The command and display variables are listed below.

Haptic Interfaces

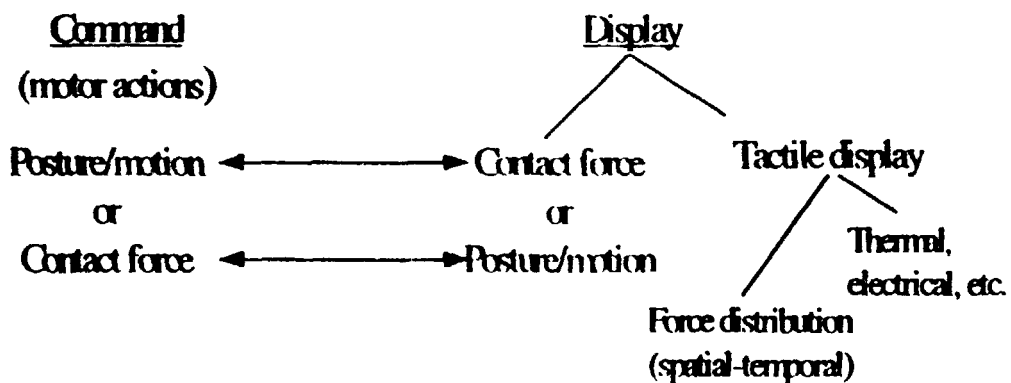
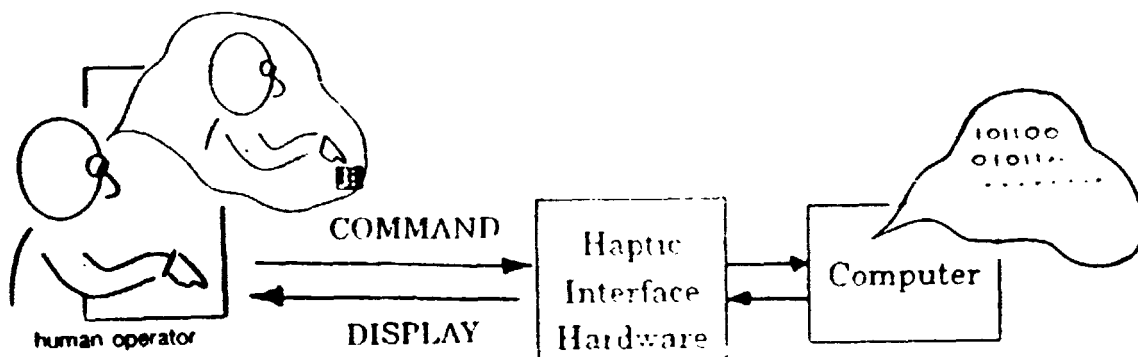


Figure 3

CLASSIFICATION OF HAPTIC INTERFACES

Haptic interfaces can be classified in several ways. First, classification is based on whether they are force-reflecting or not, as well as by what types of motions (e.g., how many degrees of freedom) and contact forces they are capable of simulating. The second type of classification is based on whether they are simulating the touch, feel, and manipulation of objects directly in contact with the skin or through a tool. A third set of important distinctions are based on whether the force display systems are ground-based, such as joysticks and other hand controllers, or body-based, such as gloves and exoskeletons.

- 1. Based on motions and/or forces (e.g., presence or absence of force reflection, degrees of freedom, types of forces)**
- 2. Ideal exoskeleton or tool handle approach**
- 3. Ground-based or body-based**

Figure 4

AVAILABLE HAPTIC INTERFACES

A variety of haptic interfaces are currently available. Computer keyboards, mice and trackballs constitute relatively simple haptic interfaces. Gloves and exoskeletons that track hand postures have more interaction capabilities and are available in the market. Although desktop and wearable force-reflecting devices have been built and implemented in research laboratories, the current capabilities of such devices are quite limited. There exist a number of examples of tactile stimulators for the finger, including pneumatic shape changers, electrocutaneous stimulators, and vibrating arrays, but none provides convincing tactile images and all are awkward to use (Ref. 2).

- **Position sensors**
- **Joysticks**
- **Point-interaction robotic devices**
- **Teleoperator masters**
- **Exoskeletal devices:**
 - **flexible (gloves and suit worn by user)**
 - **rigid links (joined linkages affixed to user)**
- **Tactile displays:**
 - **shape changers (shape memory actuators, pneumatic actuators, micro-mechanical actuators)**
 - **vibrotactile**
 - **electrotactile**

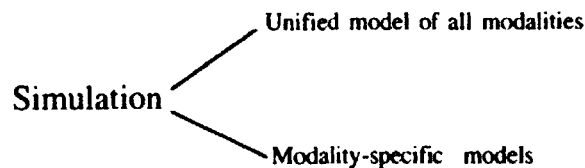
Figure 5

SOFTWARE FOR HAPTIC INTERACTIONS

Similar to the software needed to generate visual images, the software necessary to generate tactual images can be classified into three major groups: haptic interaction software, simulation of object behavior, and software for rendering tactual images. Haptic interaction software mainly consists of reading the state of the haptic interface device. Simulation of object behavior requires physical models of virtual objects. This can be accomplished either by a unified model for all the modalities (e.g., visual, haptic, acoustic) or through separate models for each modality, together with correlation algorithms for consistency among the displays corresponding to each of the modalities. The software for rendering the tactual images receives the output of the physical model and generates the commands needed to drive and control the interface device.

Software for Haptic Interactions

Device State - Reading and interpreting the state of the haptic devices.



Rendering - Control of haptic display

Figure 6

THE PHANTOM

The Phantom is a force-reflecting haptic interface recently developed at MIT (Ref. 3). It is capable of generating realistic sensations of single point contact interactions. It is essentially a robot with six degrees of freedom, capable of generating a three-dimensional force vector at its end effector, which can either be a thimble or a stylus that is manipulated by the human user.

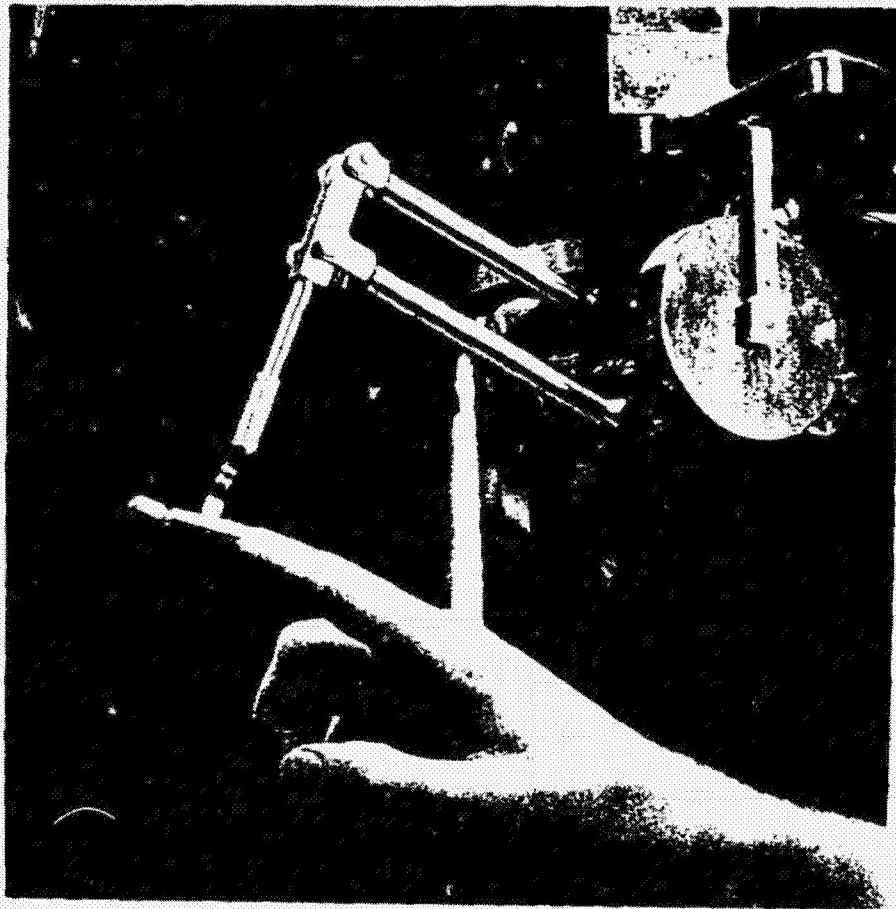
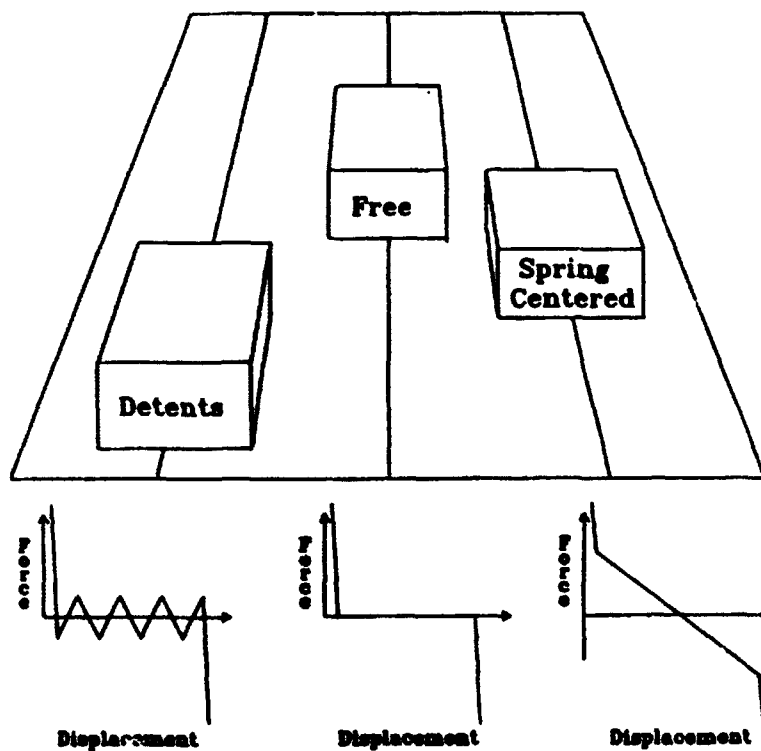


Figure 7

SLIDER SWITCHES

A variety of touch interactions have been haptically rendered using the Phantom (Ref. 4). Shown below are three types of slider switches and their respective force-displacement behavior. In addition to the properties of mass, viscosity, stiction, and surface stiffness, each cube has an underlying characteristic spring function.

Slider Switches



In addition to the properties of mass, viscosity, stiction, and surface stiffness, each cube has an underlying characteristic spring function.

Figure 8

BUTTONS WHICH "CLICK"

The feel of push buttons which "click" has also been displayed through the Phantom to the human user. Each of the buttons shown below simulate the force-displacement relationship shown schematically, but feel distinct owing to differences in the values of parameters such as stiffness.

Buttons which "Click"

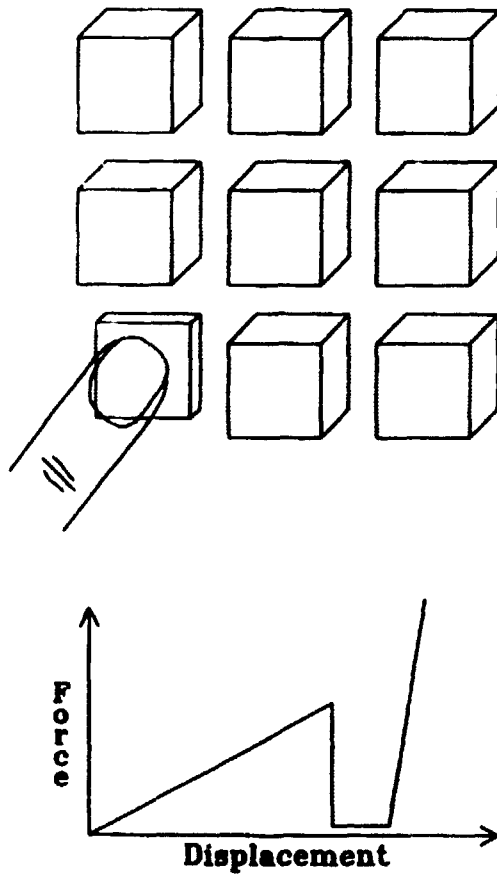
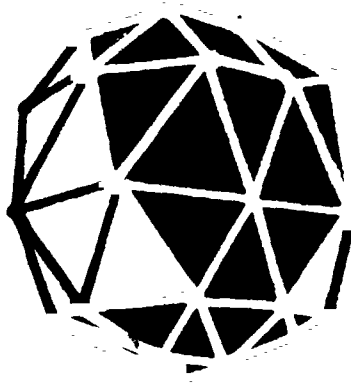


Figure 9

RENDER: RENDERS POLYGONAL SURFACES

The haptic display of arbitrarily shaped objects represented as polyhedra has been achieved with the Phantom. The rendering software uses standard graphics file formats and is capable of simulating both convex and concave surfaces.



- **Allows haptic display of arbitrarily shaped objects**
- **Uses standard graphics file formats**
- **Convex and concave**
- **Can render arbitrarily thin objects**

Figure 10

Two Phantoms have been used together to allow a user to perform two-fingered manipulation of virtual objects as shown below. Such contact interactions are analogous to the use of tools in real environments.

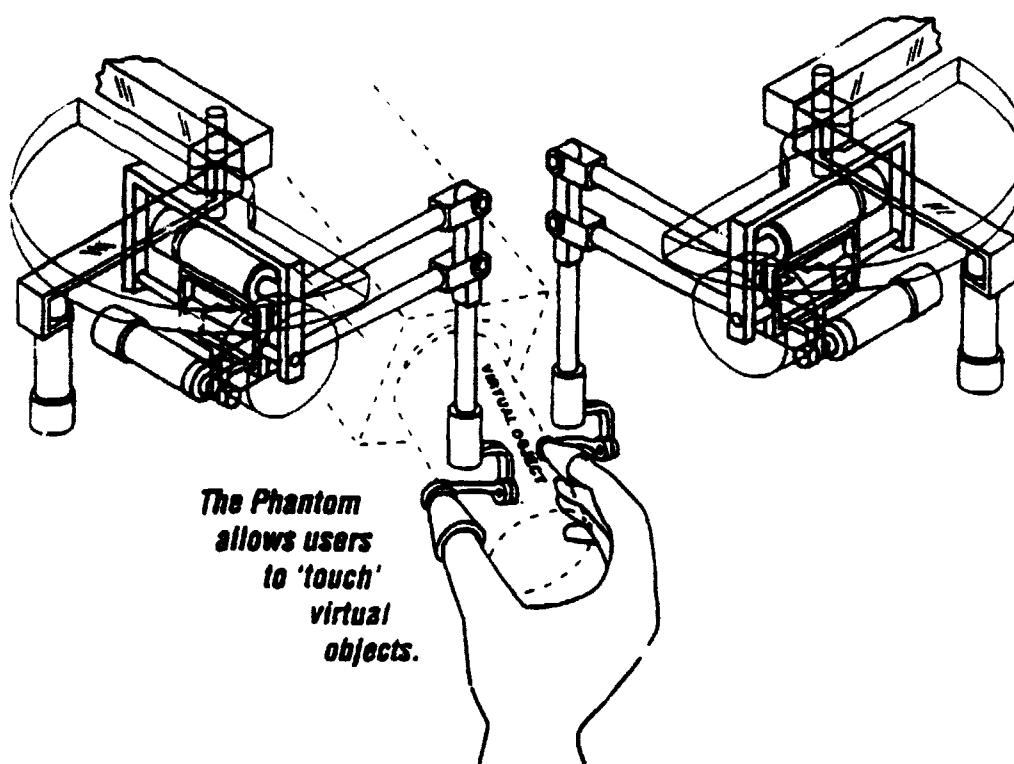
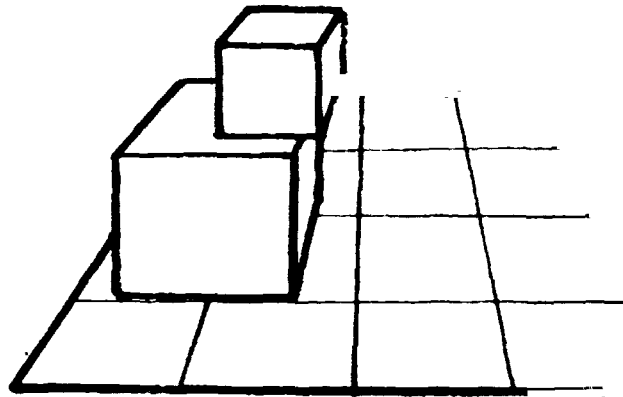


Figure 11

BLOCKS: DYNAMIC SIMULATION OF BLOCKS IN 3D

Using two Phantoms, the dynamic simulation of 3D blocks being manipulated by a user has been achieved. Both visual and haptic displays were provided to the user.



- **Two hand/finger manipulation**
- **Static friction model**
Phantom and walls, Phantom and blocks,
blocks and blocks
- **Gravity**

Figure 12

WHAT ARE ELEMENTS OF HAPTIC INTERACTION?

Listed below is a summary of the elements of haptic interactions.

- **Sensed elements:**
 - Motion, force, tactile, temperature, heat flow, current flow, pain, etc.**
- **Perceived events and states:**
 - Impact**
 - Sustained contact**
 - Slip**
 - Friction, texture**
 - Freedom/constraint in motion**
 - Compliance**
 - Curvature**
- **Workless interactions:**
 - Imposing and detecting constraint**
- **Work interactions:**
 - Force over distance, impulse, momentum and energy**

Figure 13

WHAT WE HAVE DONE

The status of the work done at the MIT AI Laboratory is given below:

- **Developed class of haptic interface
permitting force vector display - Phantom**
- **Demonstrated basic interaction elements
impact and constraint forces
object shape, motion
friction, texture
surface and object, impedance**
- **Combined basic elements to build simple
mechanical worlds:
astroid, blocks,
needle-biopsy, switch panels**
- **Developed rendering algorithms for
polyhedral objects**
- **Begun development of rendering algorithms
for visco-elastic materials**

Figure 14

HUMAN HAPTIC PERFORMANCE

A basic understanding of the biomechanical, sensorimotor, and cognitive abilities of the human haptic system is essential for further improvements in the design of hardware and software of haptic interfaces. Summarized below are some of the quantitative data (Refs. 2 and 5) on human haptic performance which set some of the design specifications of interface devices.

Tactile Sensory System:

Vibrations - detectable up to 1khz
thresholds - 0.3 to 30 microns
Spatial resolution - 1mm at fingerpad

Kinesthetic Sensory System:

Just noticeable differences (JND) for joint
angles ~ 1 to 3 degrees

Motor System:

Bandwidths: 1 to 10 Hz

Active Touch with All Three Systems

JNDs for two-fingered pinch grasp:
Length ~ 10% or less
Force ~ 7%
Compliance ~ 8%
Viscosity ~ 14%
Mass ~ 21%

Figure 15

RESEARCH NEEDS

A comprehensive program to develop a variety of haptic interfaces for virtual environments and teleoperation needs to include research in the areas shown below. Since progress in the three areas is interdependent, the desirable course of development for a challenging application is to continually build improved versions of haptic devices based on experimental data obtained from the previous version and the performance of humans, devices and the interaction between the two.

- **Human Haptics**
 - Biomechanics**
 - Psychophysics**
- **Hardware**
 - Position trackers**
 - Force displays**
 - Tactile displays**
- **Software**
 - Multimodal interactions**
 - Real time simulations**
- **Matching Human and Device Performance**
 - Comfort**
 - Simulation methods**
 - Evaluation**

Figure 16

REFERENCES

1. Durlach, N. I. and Mavor, A. S. (eds.), *Virtual Reality: Scientific and Technical Challenges*, National Academy Press, 1994.
2. Srinivasan, M. A., "Haptic Interfaces," in *Virtual Reality: Scientific and Technical Challenges*, Durlach, N. I. and Mavor, A. S. (eds.), National Academy Press, 1994.
3. Massie, T. H. and Salisbury, J. K., "The Phantom Haptic Interface: A Device for Probing Virtual Objects," in *Dynamic Systems and Control*, C. J. Radcliffe (ed.), DSC Vol. 55-1, ASME, NY, 1994, pp. 295-299.
4. Salisbury, J. K., Brock, D., Massie, T., Swarup, N. and Zilles, C., "Haptic Rendering: Programming Touch Interaction with Virtual Objects," *Proc. ACM Symposium on Interactive Three-Dimensional Graphics*, 1995.
5. Tan, H. Z., Srinivasan, M. A., Eberman, B. and Cheng, B., "Human Factors for the Design of Force-Reflecting Haptic Interfaces," in *Dynamic Systems and Control*, C. J. Radcliffe (ed.), DSC Vol. 55-1, ASME, NY, 1994, pp. 353-359.

**NEXT
DOCUMENT**

62
246

MULTI-MODAL VIRTUAL ENVIRONMENT RESEARCH AT ARMSTRONG LABORATORY

**Robert G. Eggleston
Paul M. Fitts Human Engineering Division
Crew Systems Directorate
Armstrong Laboratory
Wright-Patterson AFB, OH**

INTRODUCTION

One mission of the Paul M. Fitts Human Engineering Division is to improve the user interface for complex systems through user-centered exploratory development and research activities. In support of this goal, many current projects attempt to advance and exploit user-interface concepts made possible by so-called Virtual Reality (VR) technologies. Virtual environments may be used as a general purpose interface medium, an alternative display/control method, a data visualization and analysis tool, or a graphically based performance assessment tool. All of these uses of VR may be exploited in the development of new user-interface prototypes and supporting design tools. As a result, the Division has several active R&D efforts in these areas, as they pertain to user-machine interfaces.

The purpose of this presentation is to provide a brief overview of the range of R&D projects within the Division that involve VR technology. For the purpose of discussion, research projects are clustered into four categories:

- **Prototype Interface Hardware/Software Development**
- **Integrated Interface Concept Development**
- **Interface Design and Evaluation Tool Development**
- **User and Mission Performance Evaluation Tool Development.**

PROTOTYPE INTERFACE HARDWARE SOFTWARE DEVELOPMENT: HELMET-MOUNTED DISPLAY SYSTEMS

The Division has been a pioneer in the development of visually coupled systems (VCS) technology which centers around the design of helmet-mounted display (HMD) devices. Beginning in 1966, a line of HMDs has been produced, ranging from small field-of-view (FOV) monocular systems to wide FOV stereoscopic ones. R&D activities address the development of high luminance miniature CRTs, relay optical systems, head gear, militarized image generators, trackers, and all technological and engineering aspects required to produce fully integrated systems for use in airborne platforms. The technical advancements made in this work over the years have provided the foundation from which the virtual reality industry has emerged. For example, the need for a method to reliably track head position without restricting operator movement or emergency egress stimulated the development of the magnetic tracking concept now used in many VR systems. The first successful implementation of such a device was sponsored by the Division and achieved in collaboration with our staff.

Current VCS development efforts are concentrated on:

- standardized, plug compatible VCS systems for integration into military platforms
- color CRT that can operate in high (luminance) ambient environments
- system and component performance measurement methods and standards
- psychophysical performance assessment

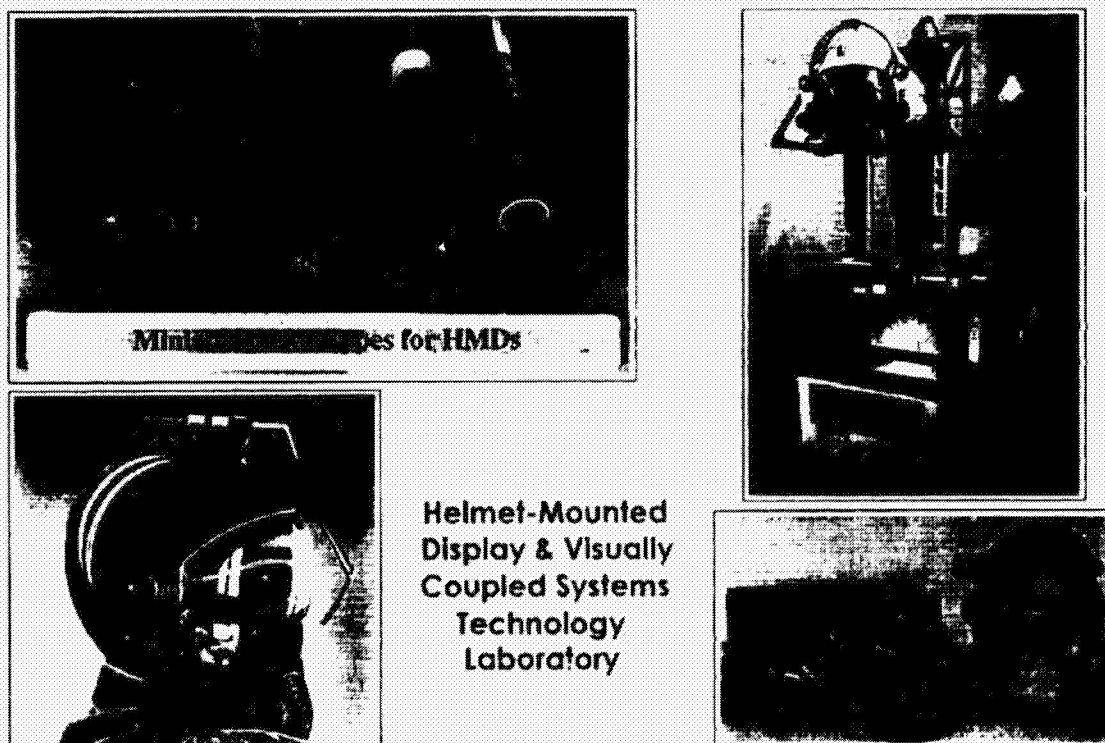


Figure 1

VCS CONTROL/DISPLAY CONCEPTS

In addition to developing the hardware and software technology infrastructure for the VCS system, we also develop advanced interface concepts around the technology. One example is the design of an interface device that supports head aiming (to select and track targets) and head-referenced instrument displays (flight and weapons) that are always available to the operator. This work involves the design of display formats, symbology, and control methods. It also includes a wide range of human performance research to insure compatibility of the concept to both the user and the expected real world task conditions. Several concepts have been developed and evaluated in simulated and actual flight tests.

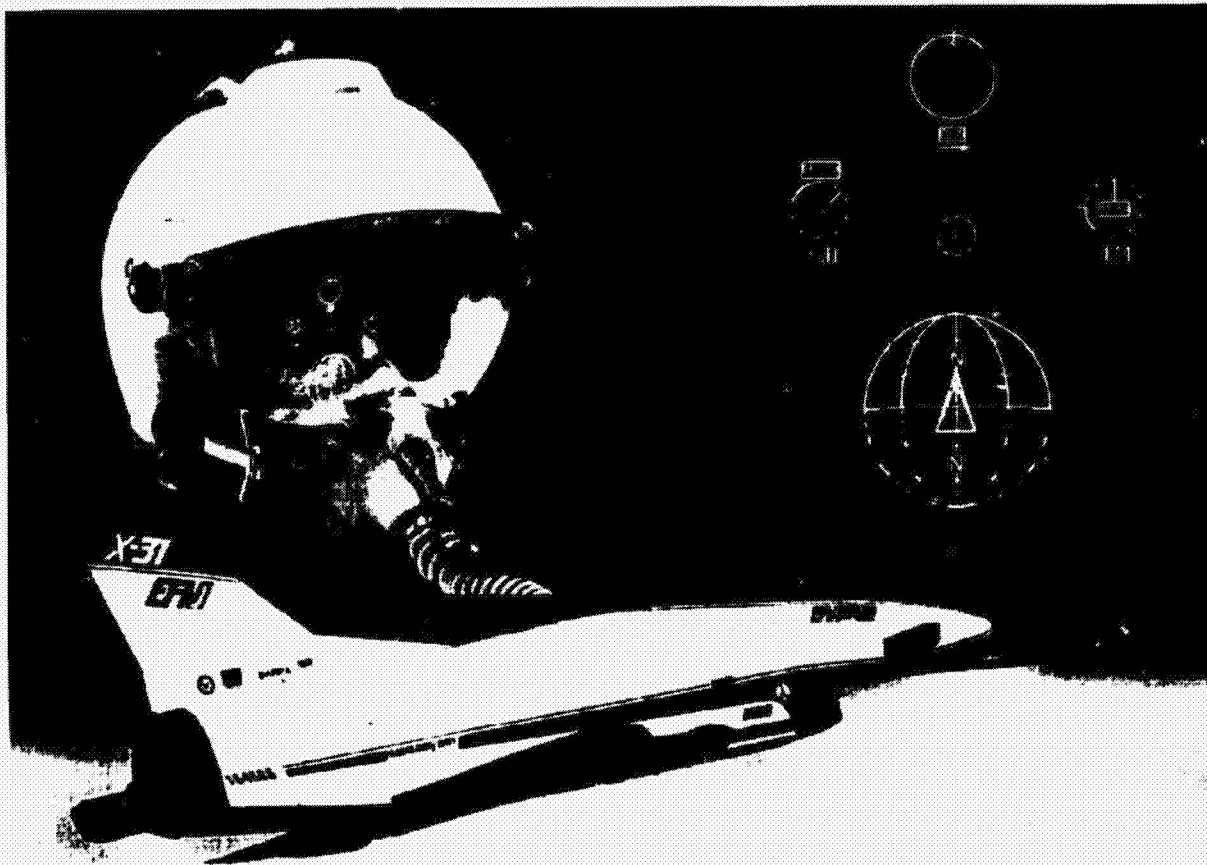


Figure 2

THREE-DIMENSIONAL AUDIO DISPLAYS

Similar to the VCS R&D efforts, Armstrong Laboratory has also been a leader in the development of 3-D audio systems, including hardware concepts, psychophysical analysis, application systems, and flight testing. Our first localized display was designed and fabricated in-house in the late 1980's. New systems compatible with operational avionics systems are being developed and evaluated.

The geodesic dome shown below is housed in an anechoic chamber located in the laboratory. This is the only available facility in the country that can support the full range of technical studies needed to advance this technology.

A study was recently completed that addressed target acquisition performance with and without the aid of a 3-D sound localizer. A summary of the results of that study is shown in the next two charts. Each graph depicts target search times as color zones (shown here as shapes of gray) over the search space defined in terms of azimuth and elevation angle. The first graph displays times when search was completed visually without the aid of localized sound cueing. The next graph shows the search time map when localized sound cueing was used. The improvement with sound cueing is obvious from visual inspection of the graphs.

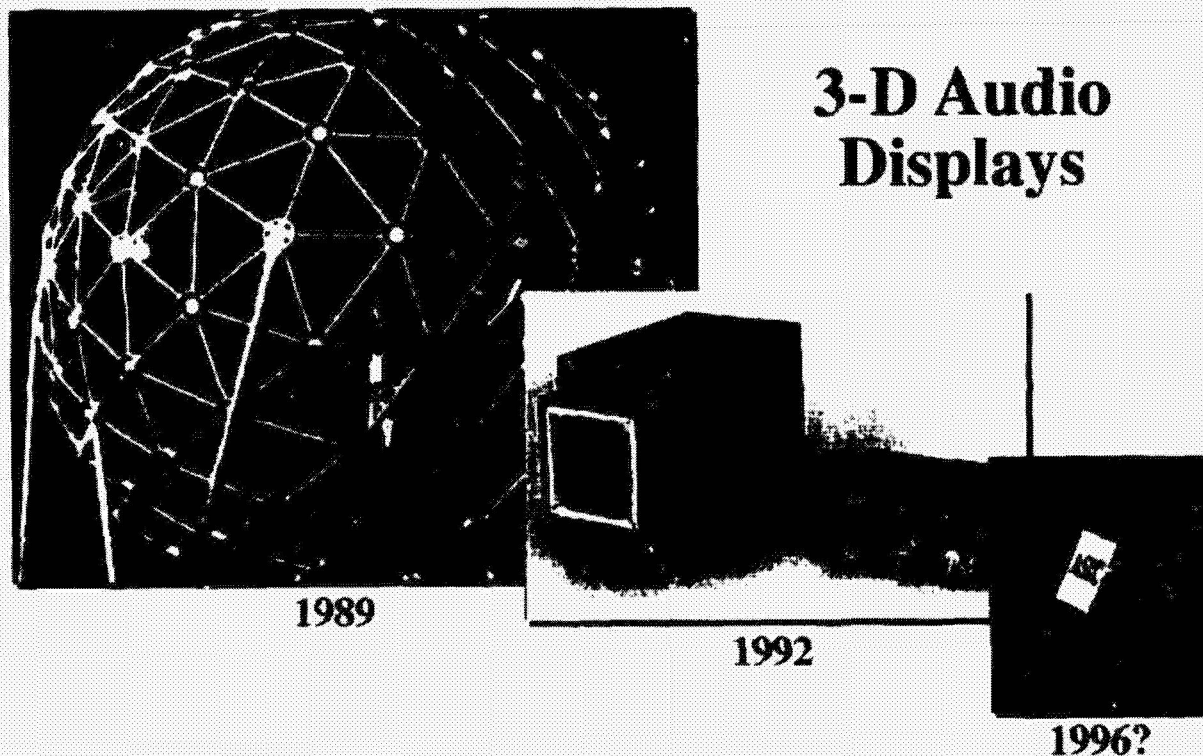


Figure 3

3-D AUDIO DISPLAYS

Visual Only Search

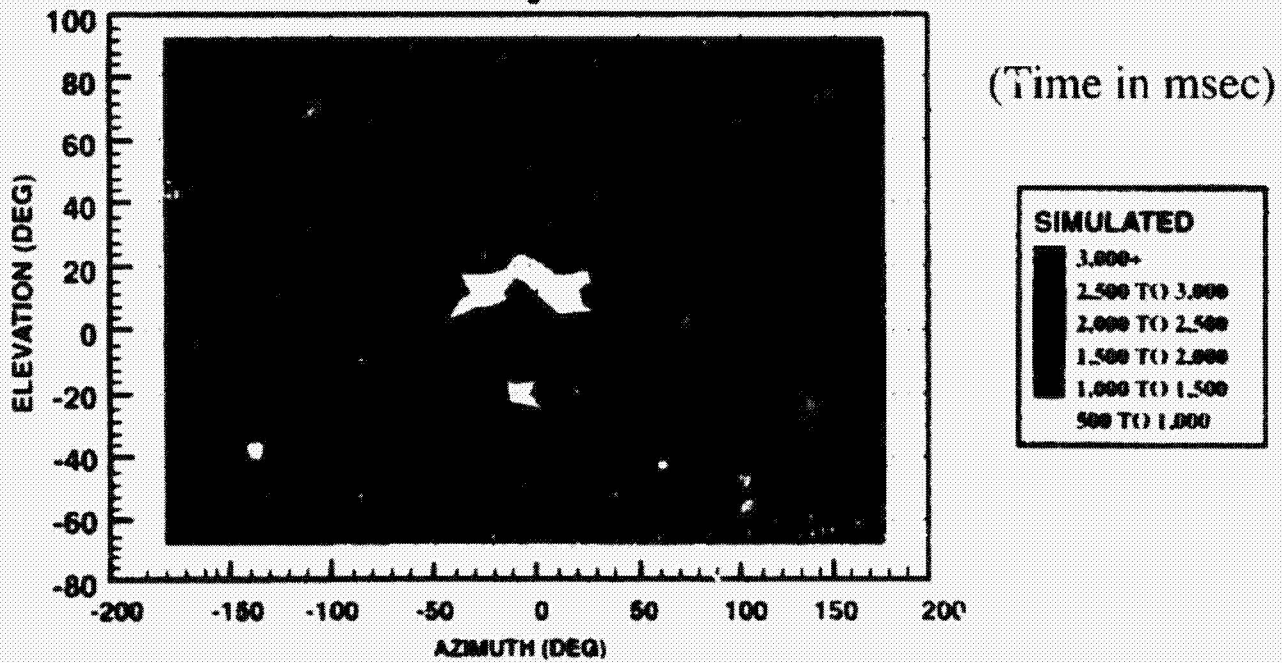


Figure 4

Combined 3-D Audio/Visual Search

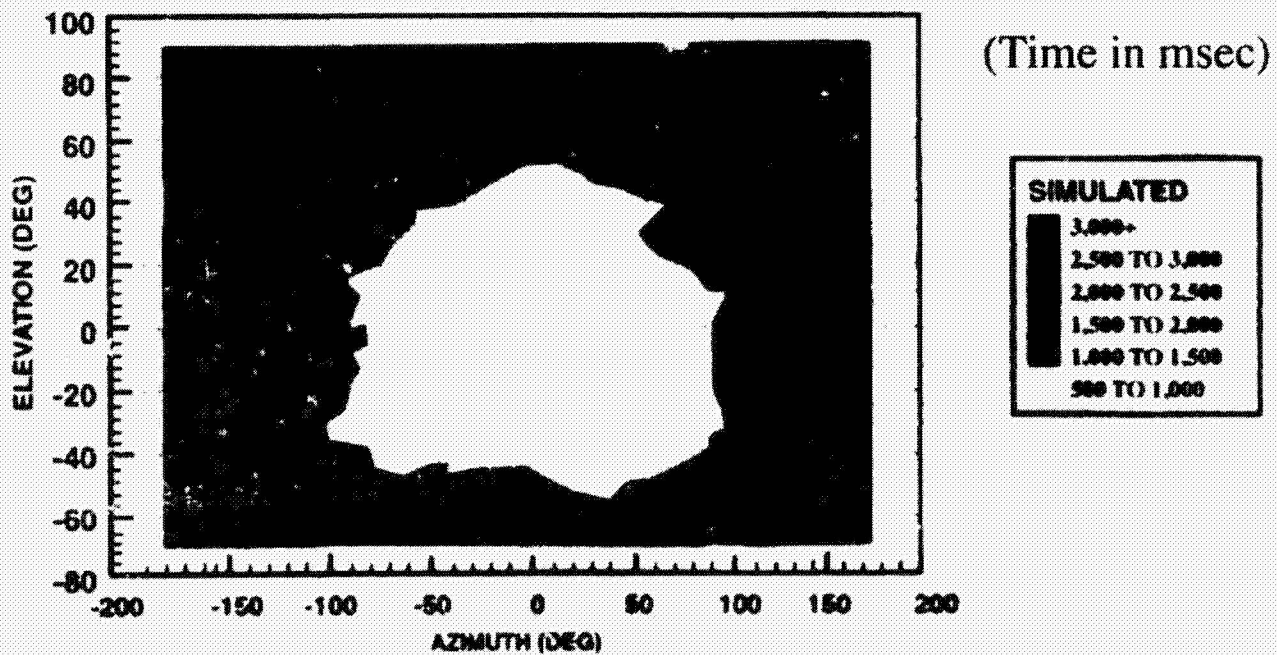


Figure 5

ALTERNATIVE CONTROL TECHNOLOGIES

Several projects in the Division are involved with developing and exploring the feasibility of novel interface methods as alternative I/O devices for use in a wide variety of application domains. These include the development of a gesture recognition system that combines a data glove which senses figure postures and gestures (coupled with a position/orientation tracker) with a neural net classifier. A demonstration system is being constructed that recognizes American sign language. The device may be used as an input mechanism to a system/computer or as a means to support nonverbal communication between people. The Veterans Administration has expressed interest in this work as an assistive technology.

A facial sensor testbed has been developed to determine if continuous speech recognition accuracy can be improved by "reading" facial expressions during speech. The basic idea is to use facial patterns as a basis for disambiguating classification uncertainties found with current hidden Markov model speech recognition systems.

Perhaps the most unusual alternative control technology under study is called Brain Actuated Control (BAC). As the name suggests, we are exploring ways brain states can be detected and used as a control decision which is then executed by the use of standard servomechanism principles in a closed-loop system. We have demonstrated reliable two-state based brain control, with the operator voluntarily enhancing and suppressing the magnitude of a certain EEG frequency. Continuous control has also been demonstrated through exploitation of a composite voluntarily produced EEG-EMG signal. This is another area where our work has spin-off potential for use with disadvantaged individuals.

Facial Sensor Test Bed



Hand Gesture Recognition

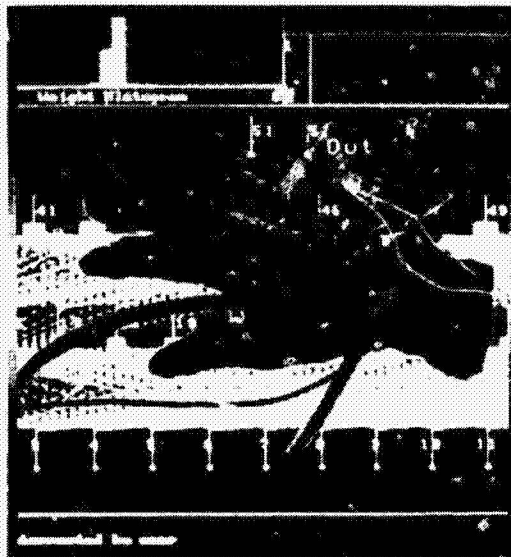


Figure 6

Brain Actuated Flight Simulator

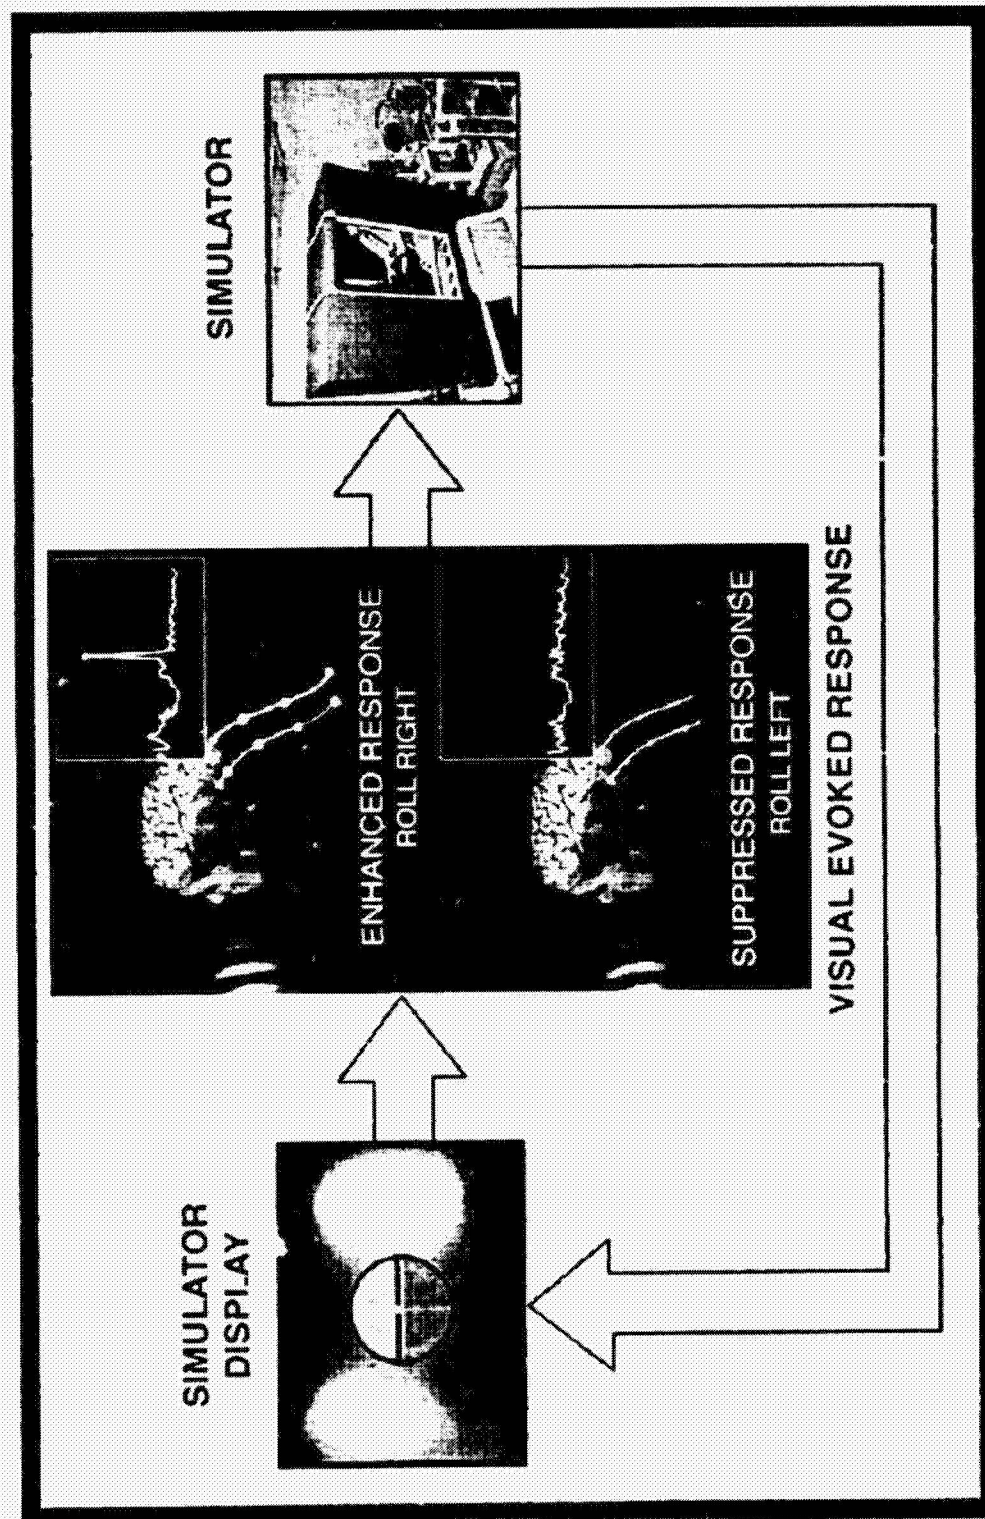


Figure 7

INTEGRATED INTERFACE CONCEPT DEVELOPMENT

Beginning with the Visually Coupled Airborne Simulation System (VCASS), which became operational in the early 1980's, we have had the ability to produce virtual interfaces or virtual cockpit concepts for investigation. This system probably marked the beginning of the new era of VCS technology that has become known as Virtual Reality, Virtual Environment, or Synthetic Environment technology. VCASS provides a 120 degree (horizontal) by 60 degree (vertical) instantaneous FOV in a head-mounted stereoscopic display system. We have coupled it over the years with data gloves, a simple force reflecting device, and 3-D audio to produce various forms of virtual cockpits. Unfortunately, limitations in computer graphics and system time delays have interfered with our ability to adequately assess the value of these innovative interfaces.

We have recently opened a new facility called the Synthesized Immersion Research Environment (SIRE), that provides a full capability immersive VR system. It integrates haptic displays with 3-D sound and VCS technology and also includes a large screen projection system. SIRE supports a wide range of human performance research, including investigations of augmented cockpit concepts which couple VR concepts and alternative control concepts with more conventional panel mounted displays.

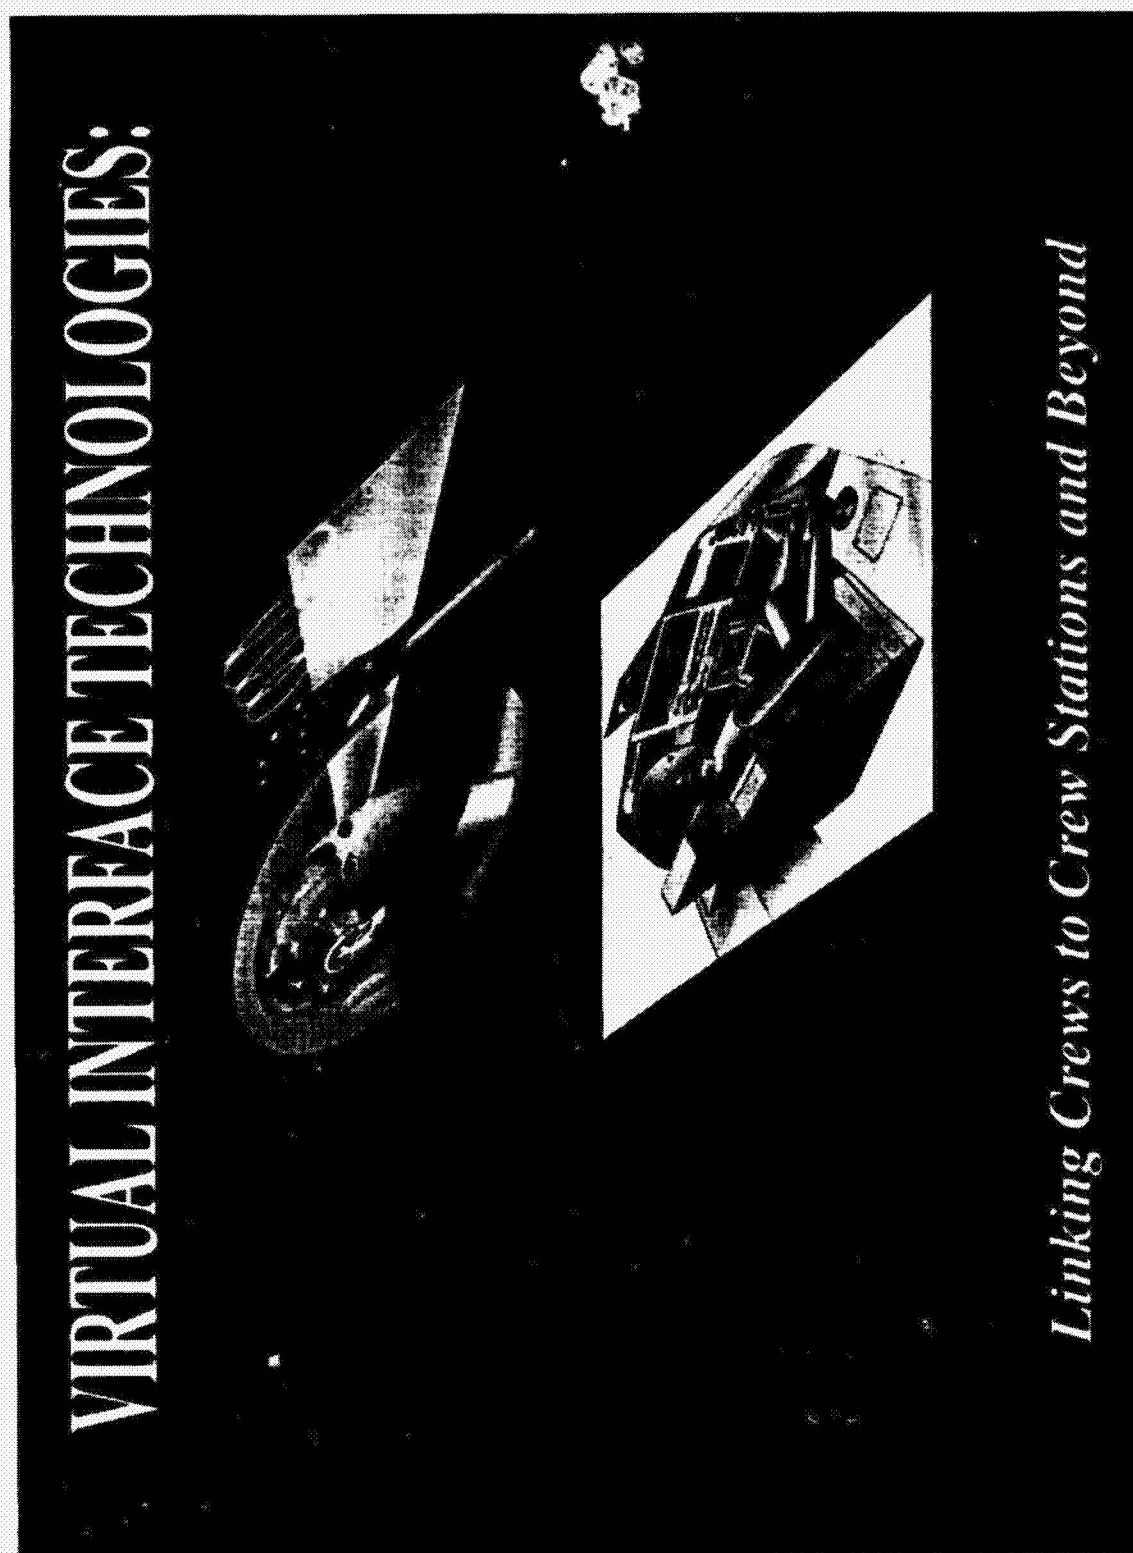


Figure 8

CONTROL INTEGRATION

INTEGRATED INTERFACE CONCEPT DEVELOPMENT

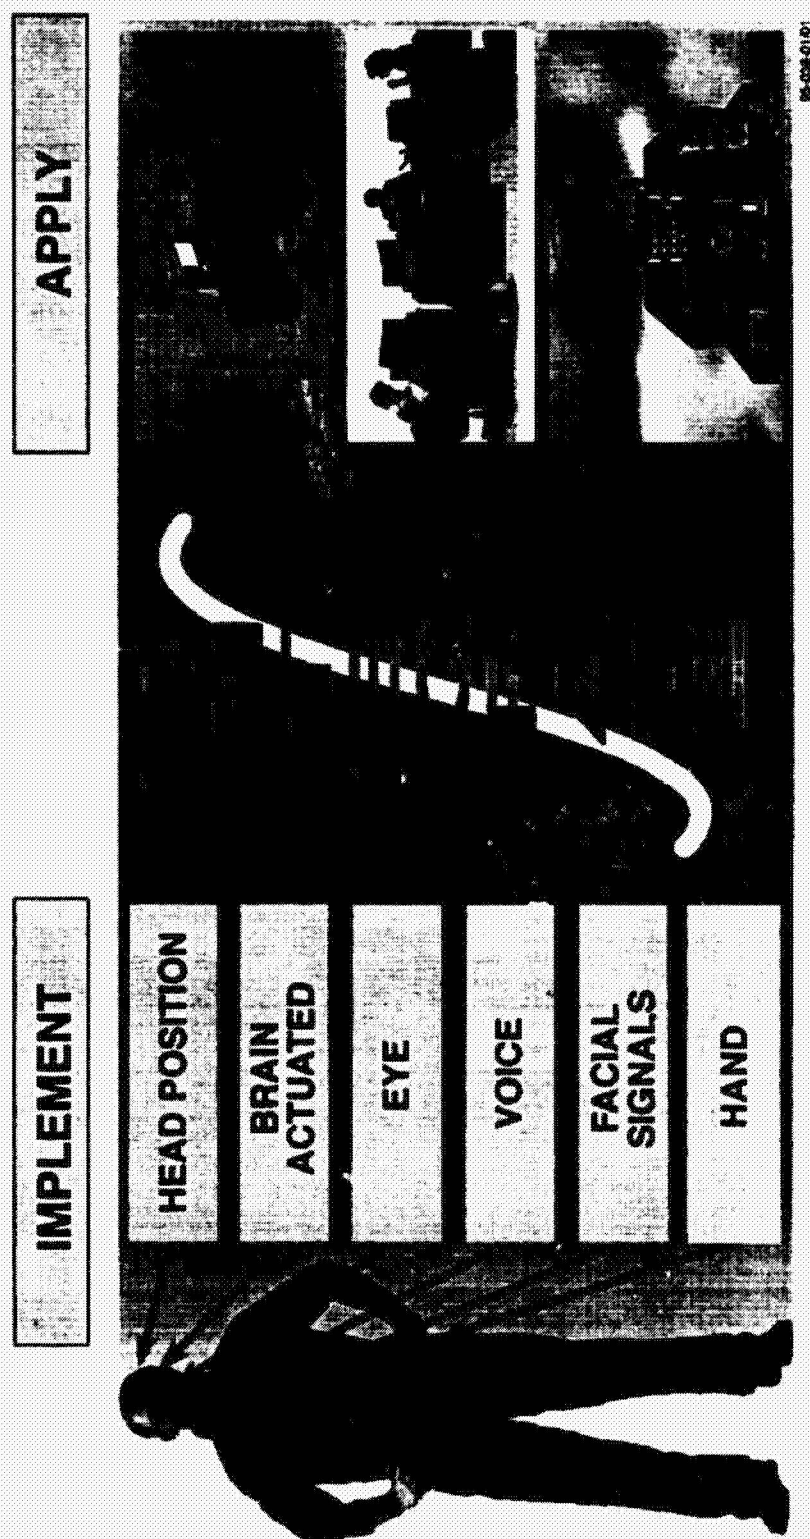


Figure 9

INTEGRATED INTERFACE CONCEPT DEVELOPMENT



Figure 10

INTERFACE DESIGN AND EVALUATION TOOL DEVELOPMENT

Historically, it has often proven difficult to transition user-centered interface ideas and concepts into operational systems. There are many reasons for this breakdown in the development process. Most agree that one reason is the lack of adequate interface design and evaluation tools that fit properly into the systems engineering process used in advanced design of large-scale systems. In an effort to overcome this difficulty, we have been developing tools and concepts to support the design and evaluation process.

One area of research important to this goal is engineering anthropometry. In this area, we perform studies that allow us to develop relevant anthropometric data bases. For example, we assess the ability of different body types (size, strength, etc.) to perform different actions often required in maintenance tasks. Data bases like these have been combined with computer graphics systems to produce interactive visualization and analysis tools for assessing system designs (computer representations). VR has recently been used to allow the designer to assume different body sizes during the analysis.

INTERFACE DESIGN AND EVALUATION TOOL DEVELOPMENT

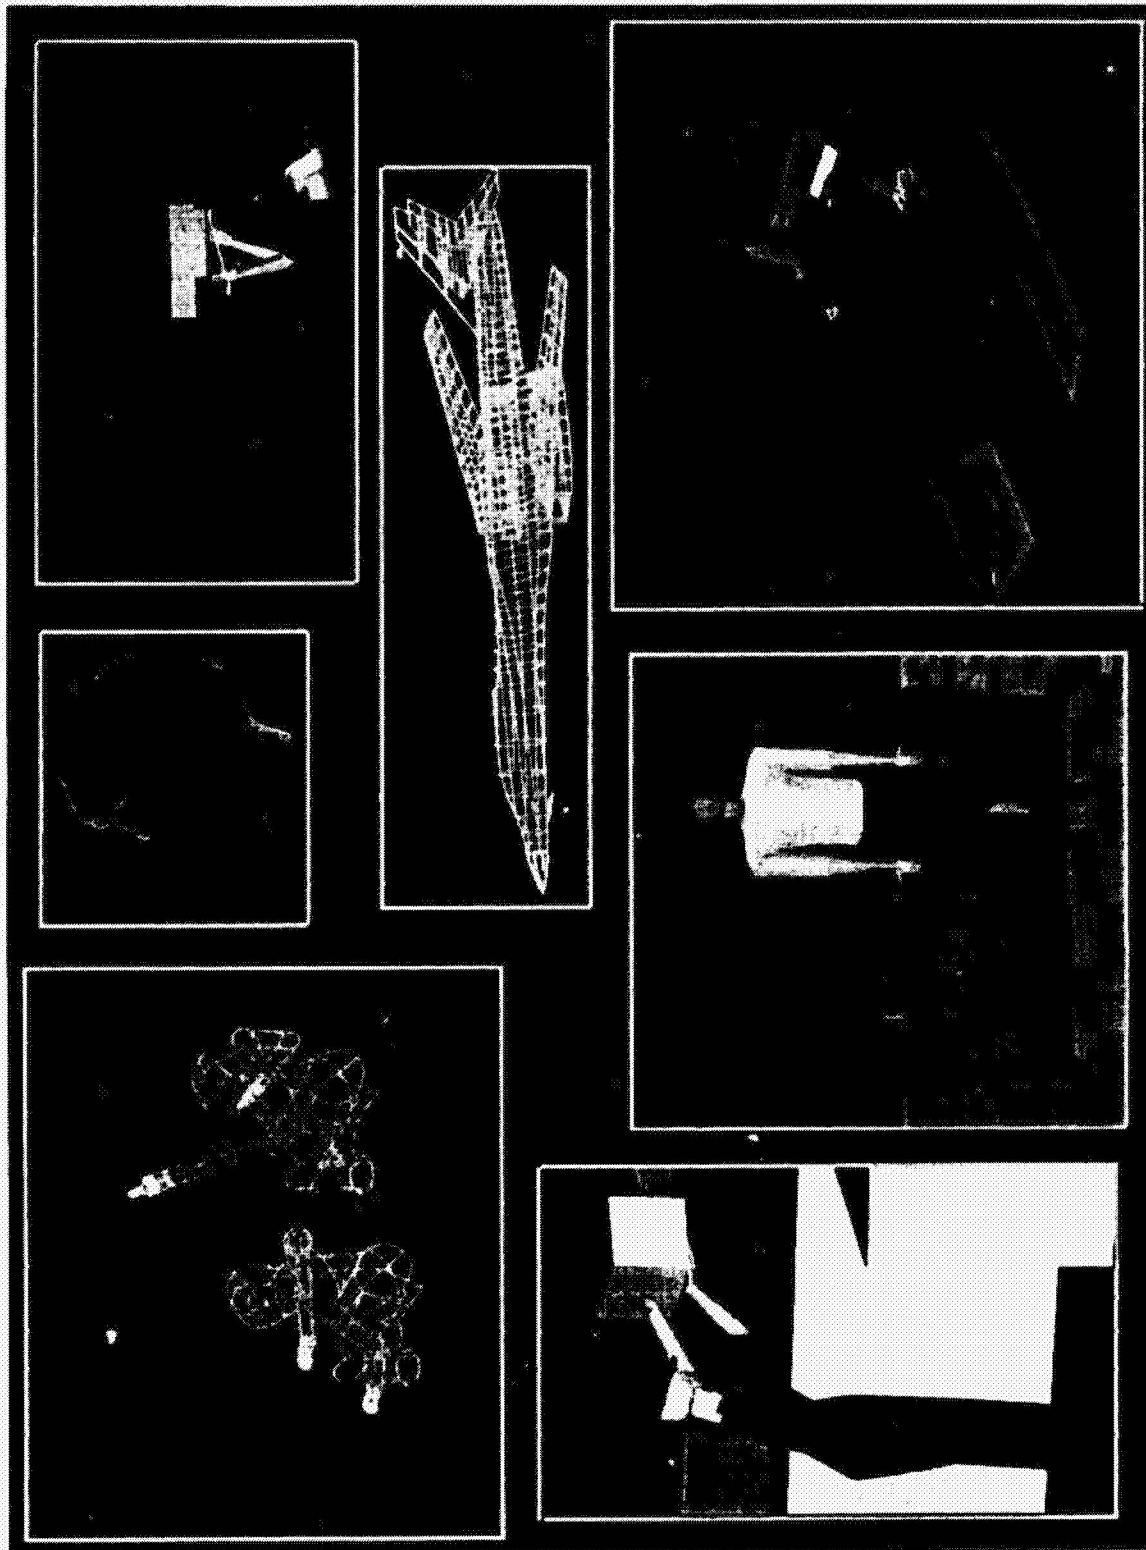


Figure 11

INTERFACE DESIGN AND EVALUATION TOOL DEVELOPMENT

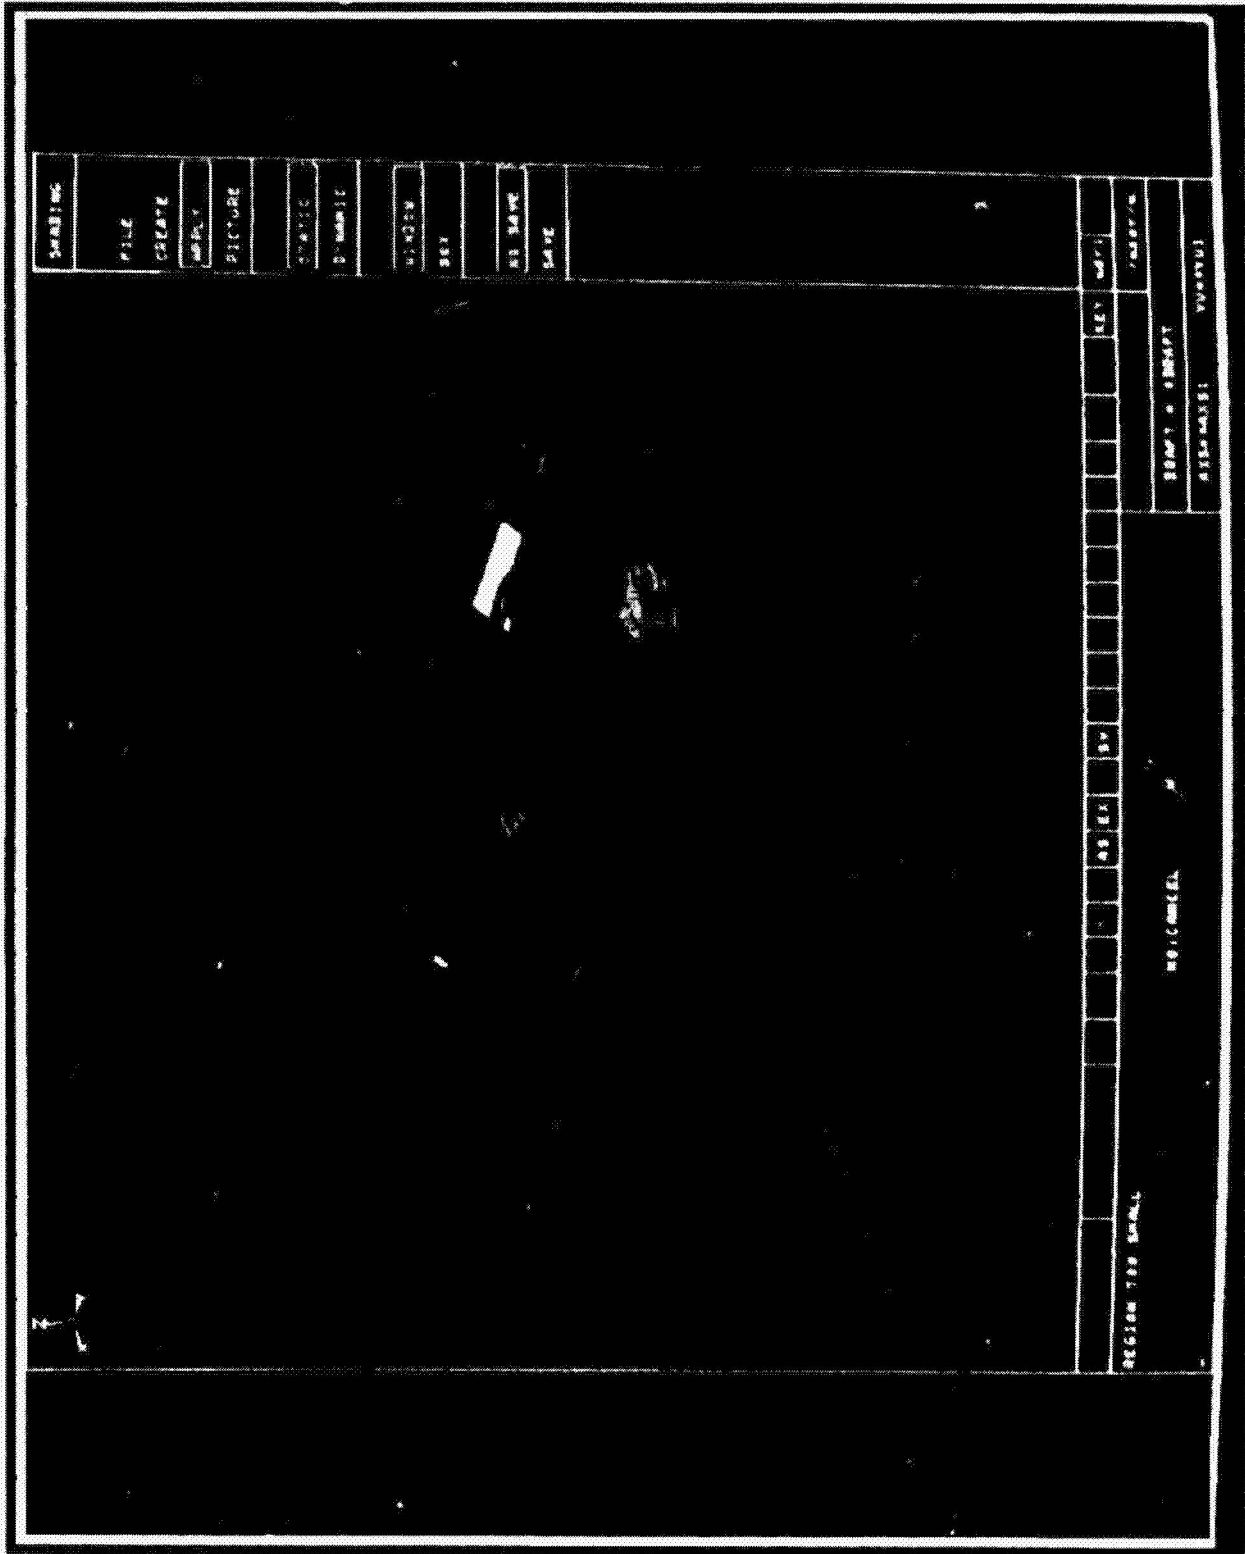


Figure 12

INTERFACE DESIGN AND EVALUATION TOOL DEVELOPMENT

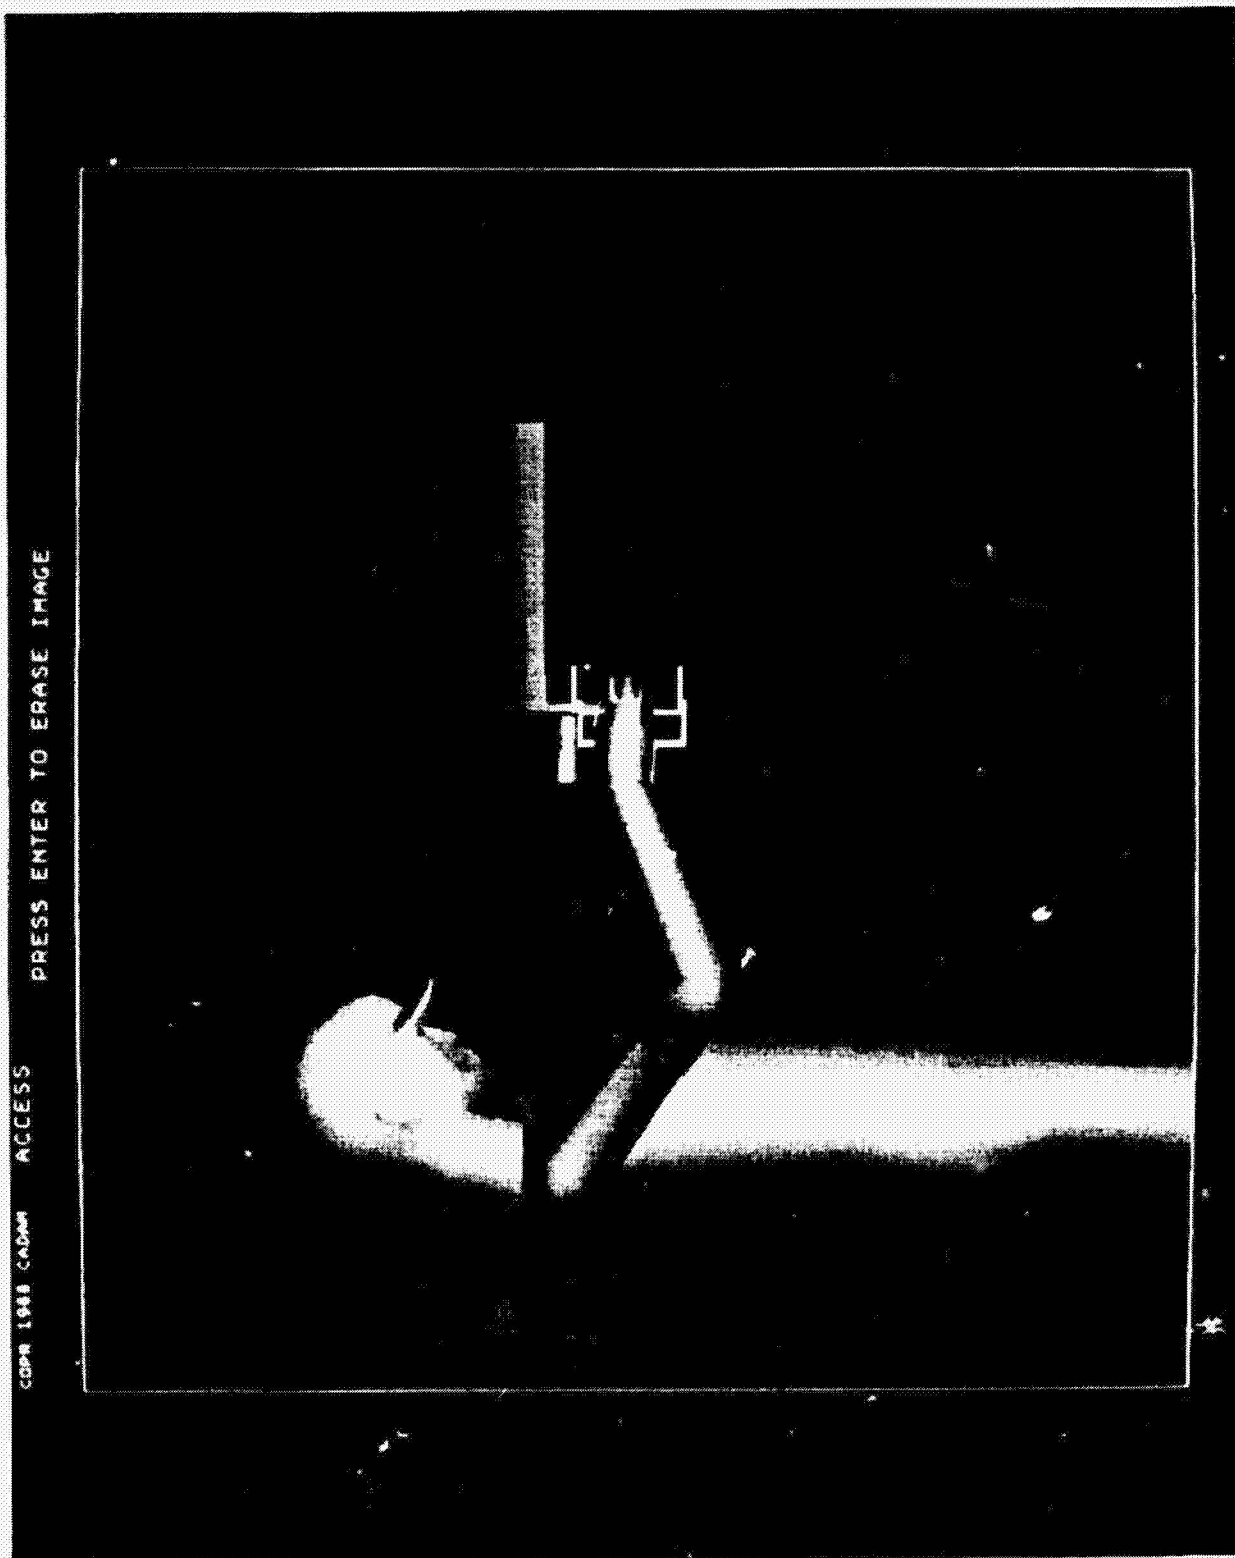


Figure 13

HUMAN CENTERED DESIGN ON THE CAD/CAE PIPELINE

Another effort has concentrated on the development of active information using multi-media technology. The goal is to be able to place human performance data into the CAD/CAE environment used in design. This research has addressed ways to improve access to information, ways to help the designer understand what questions they need to ask, and ways to facilitate opportunistic searches. In addition, we have invested in the production of a new tool called a perception and performance prototyper. This software product allows the designer to transition from equations and graphical presentations and analysis to experiential methods that allow the designer to experience the perceptual and performance consequences of their design decisions.



Figure 14

HUMAN CENTERED DESIGN ON THE CAD/CAE PIPELINE

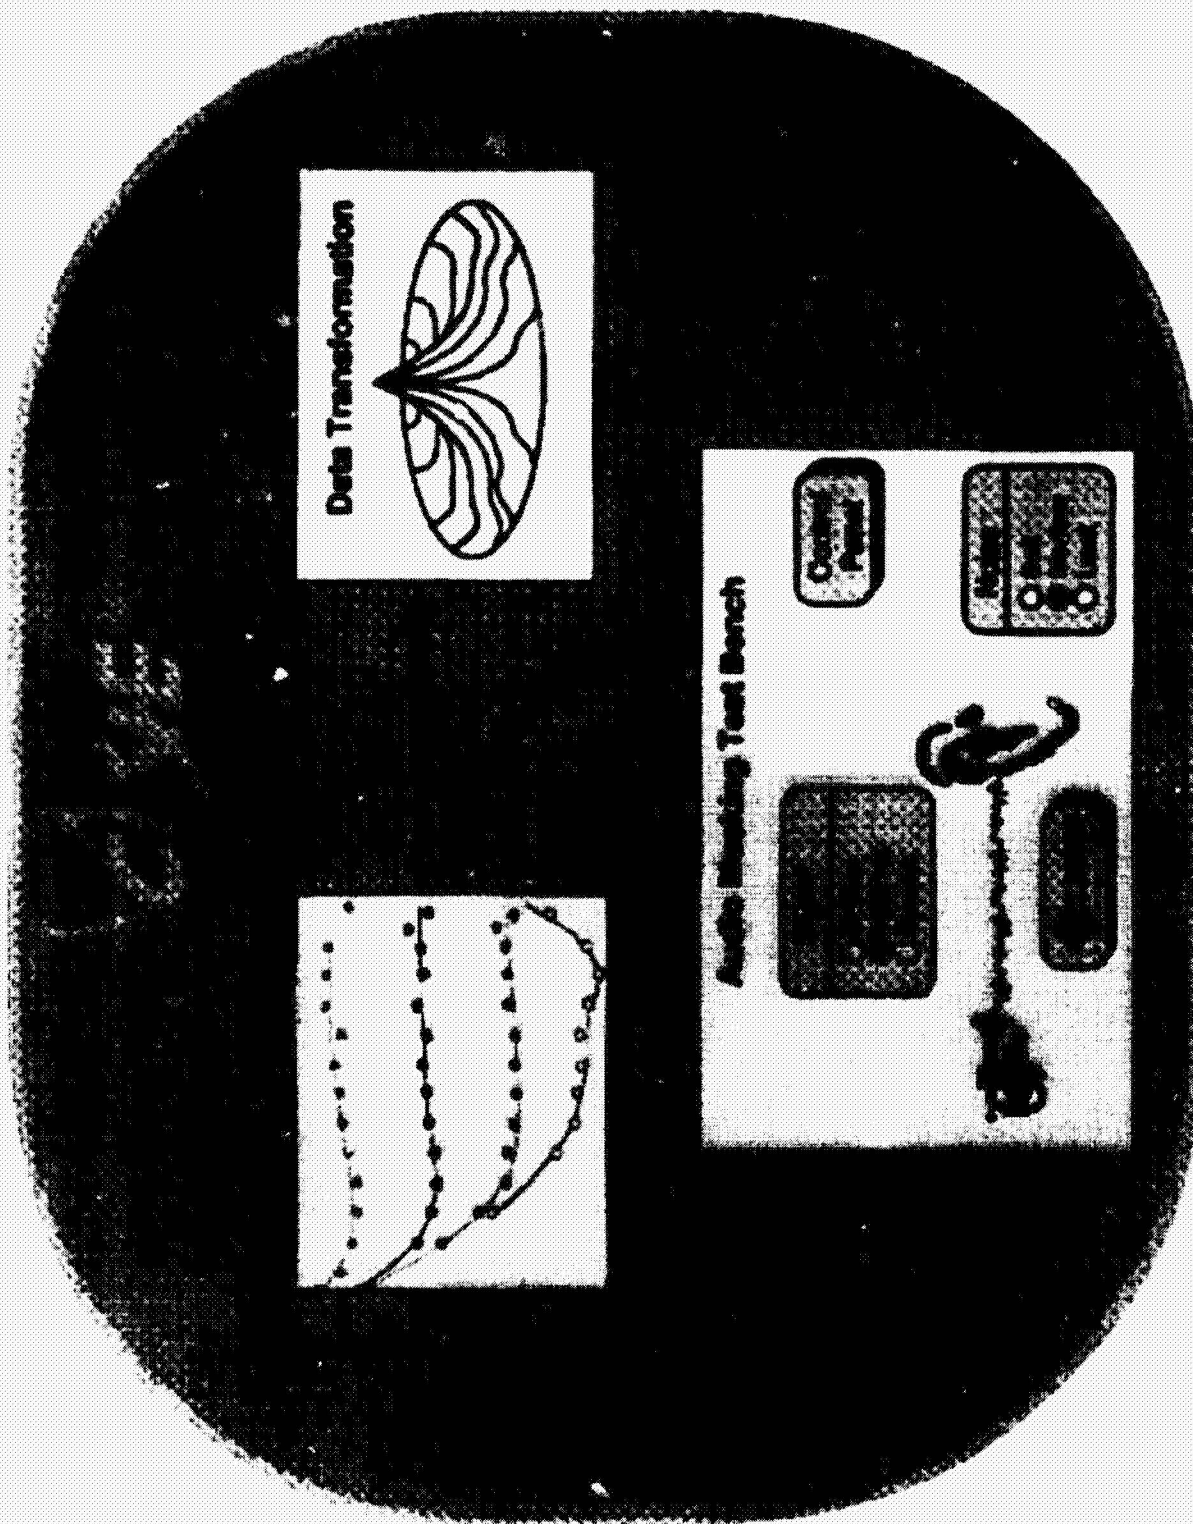


Figure 15

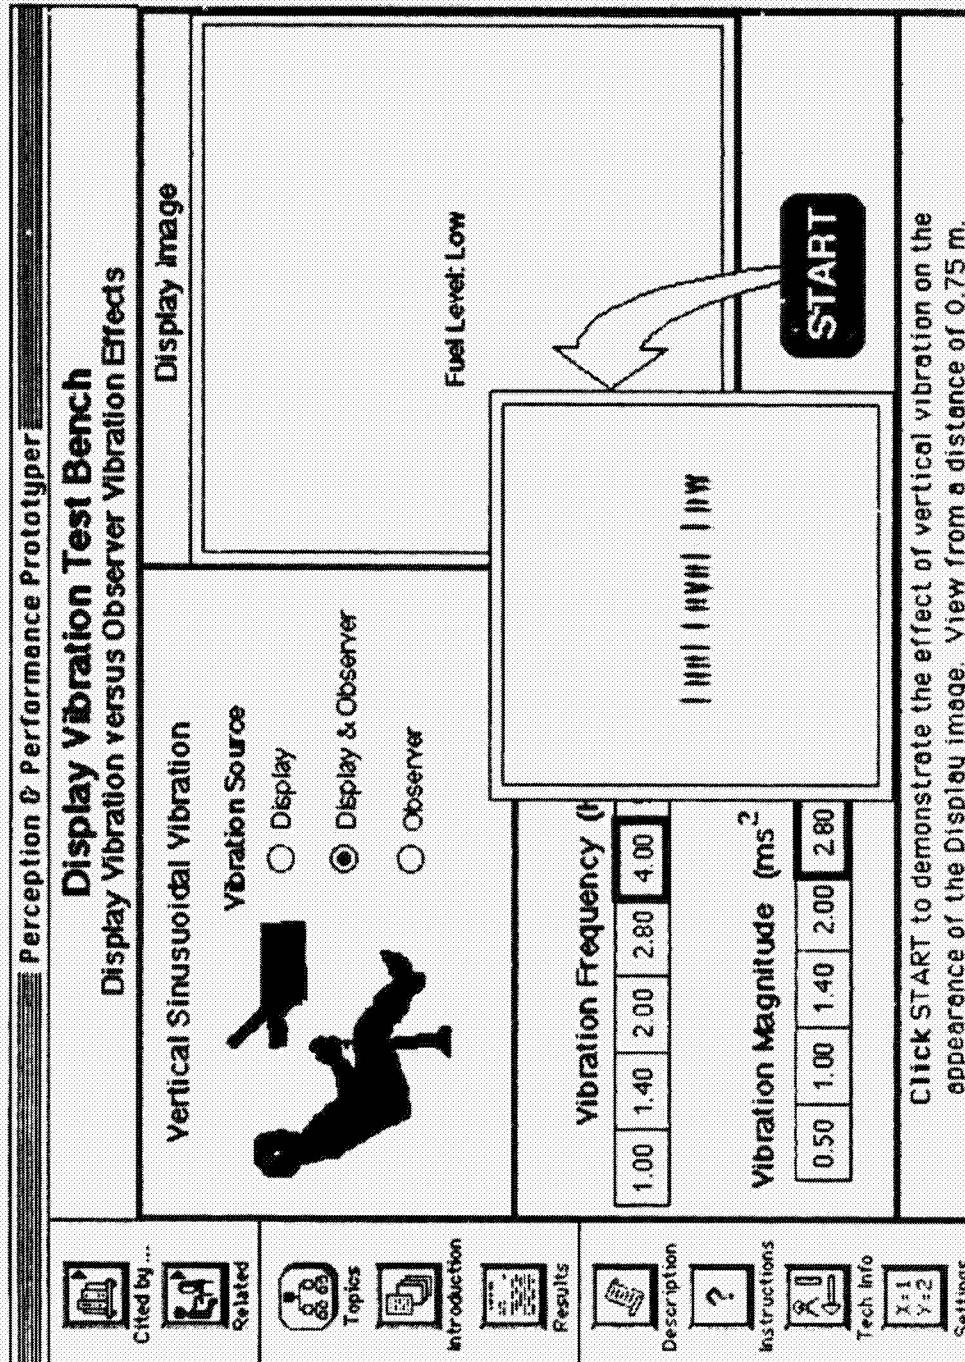


Figure 16

USER PERFORMANCE AND VIRTUAL SYSTEMS

It is clear that the potential applications for Virtual Technology are enormous. DoD, for example, is moving ahead with the use of this technology as both a tool and an end product in the development of advanced systems. Distributive Interactive Simulation (DIS) is being considered as a test bed to investigate advanced concepts under operational like conditions. Some of these systems may be virtual; others physical. In this and all other work involving VR technology, we suffer today from a lack of adequate understanding of how user performance interacts with properties of the VR technology itself. Without information like this we will not be able to determine to what degree user performance in, say, a DIS environment, reflects what can be expected in the real world, and what is due to characteristics of the VR simulations. Information of this nature is also needed to aid in design decision making, and to provide performance-based bench marks or requirements to guide VR technologists in advancing the state of the art.

We have undertaken a research project that focuses on this issue. One goal of this effort is to establish a quantitative relation between VR system properties and user performance. Our plan is to do this with a series of standardized benchmark tasks that cover a wide range of perception, perceptual-motor, and cognitive activities. To date we have completed studies that demonstrate the cost of VR system time delay on manual tracking and control and aimed movement control. The effects of several VR system variables on size-distance judgments in a virtual environment have also been studied.

The next chart depicts the fact that user performance based design data requires a link to be made between properties of user performance with properties of VR systems. The user performance properties must, in turn, be linked to specific application domains in order to map user performance to task performance. One goal of our research is to provide the desired user-performance data and the tools needed to complete the mapping to a wide variety of task domains.

Figures 18 and 19 provide a brief summary of the results from some experiments that investigated simple aimed movement performance which was accomplished in a virtual environment (VE). In general, the data follow Fitts Law over an ID range of 2 to 6, but show departure from linearity around an ID of 7. In comparison with physical aimed movement, performance by highly practiced subjects was degraded by about 35-50% in a VE (see Fig. 19). Further, the effect of time delay in the VR system caused movement times to increase substantially under ID 5 conditions, somewhat for ID 6 conditions, and very little at other ID levels (see Fig. 19). Data like these provide the basis for constructing design trade-off nomographs like the example shown in Fig. 20.

A design trade-off nomograph (see Fig. 20) is a form of data representation that highlights the relation between a measure of user performance and two or more measures of VR system properties. This form of representation allows a VR system designer to easily see how user performance interacts with system properties. It can be used for defining system requirements and for performing trade-off analyses to select between alternative design options that can meet a requirement. Thus, it aids the designer in seeing the connection between user performance in a way that is clearly connected to system properties, but at the same time allows the designer to bring in technology risk, cost, and schedule factors in the process of deciding how to meet a performance goal.

USER PERFORMANCE AND VIRTUAL SYSTEMS



Figure 17

Mean Movement Time by ID and Study

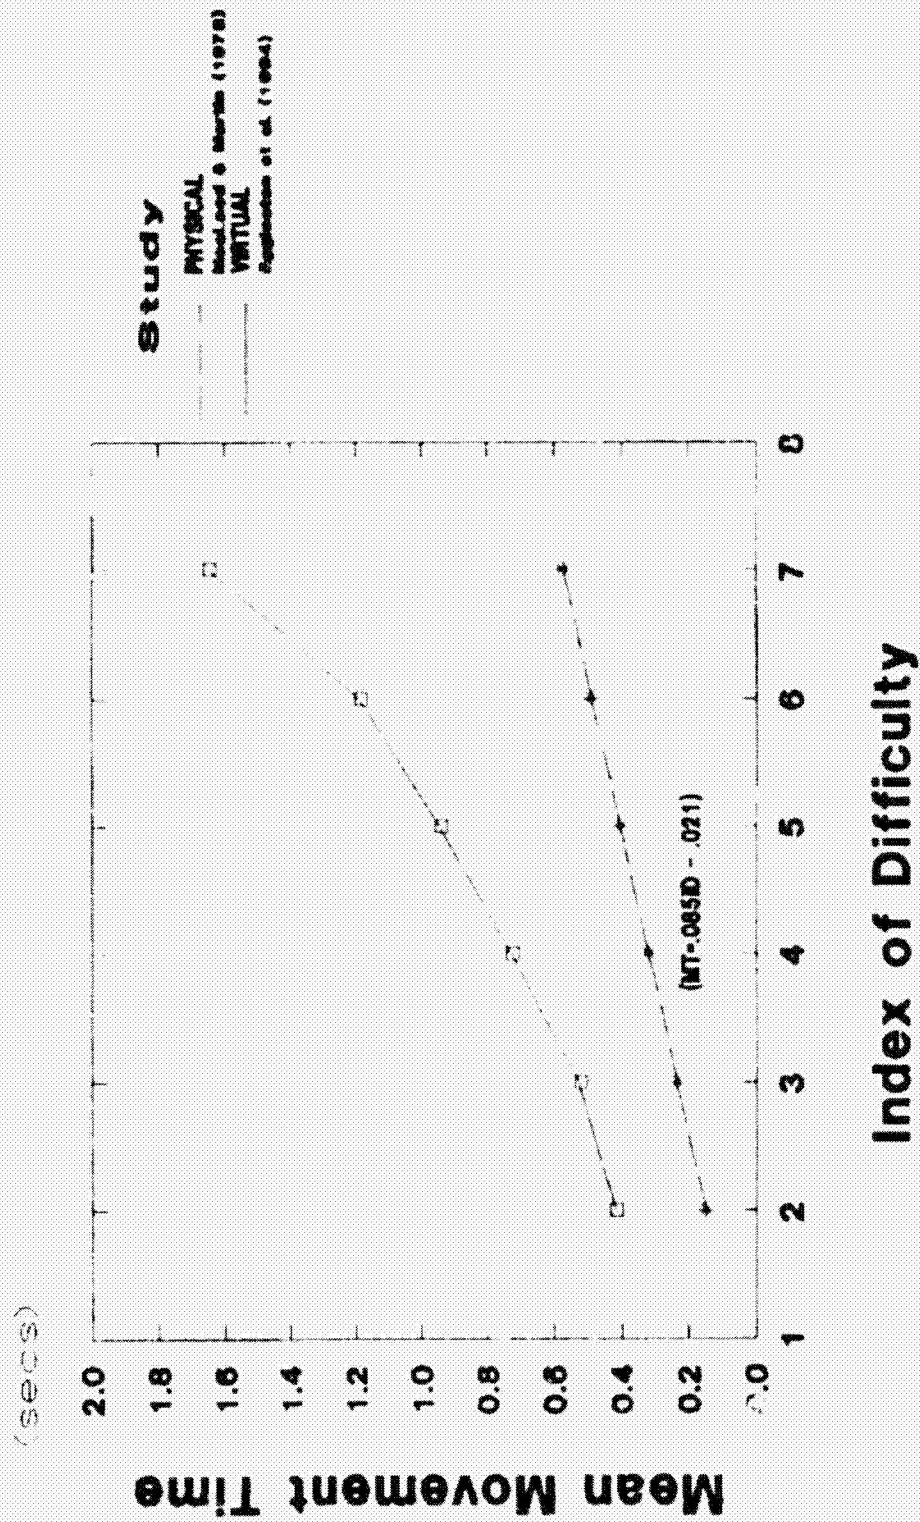


Figure 18

Reciprocal Tapping Task

Mean Movement Time by ID and Study

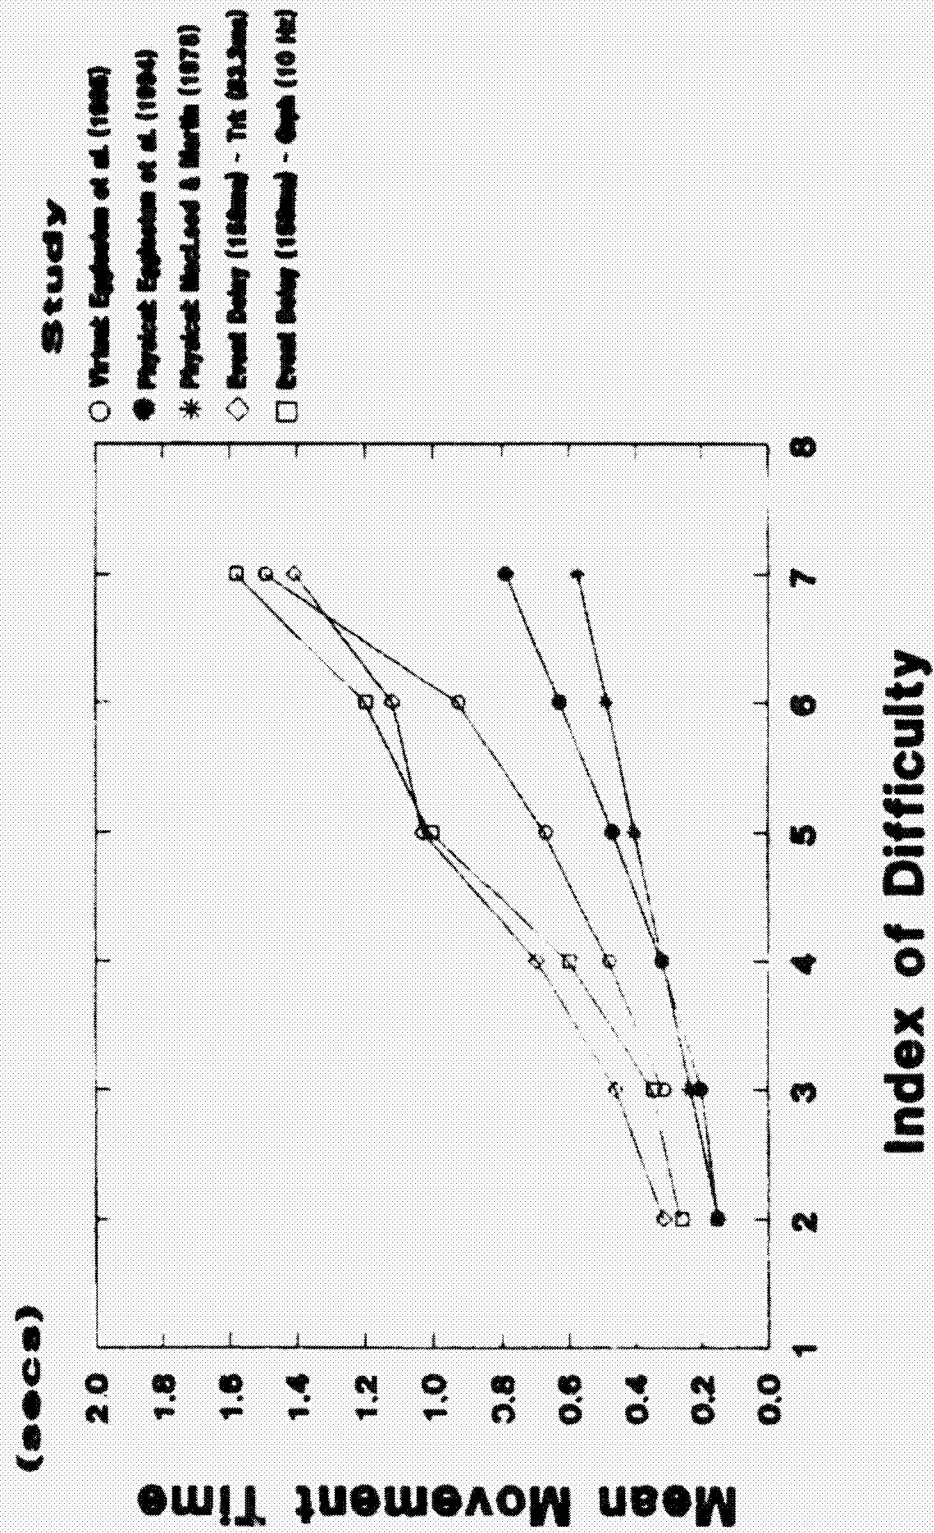


Figure 19

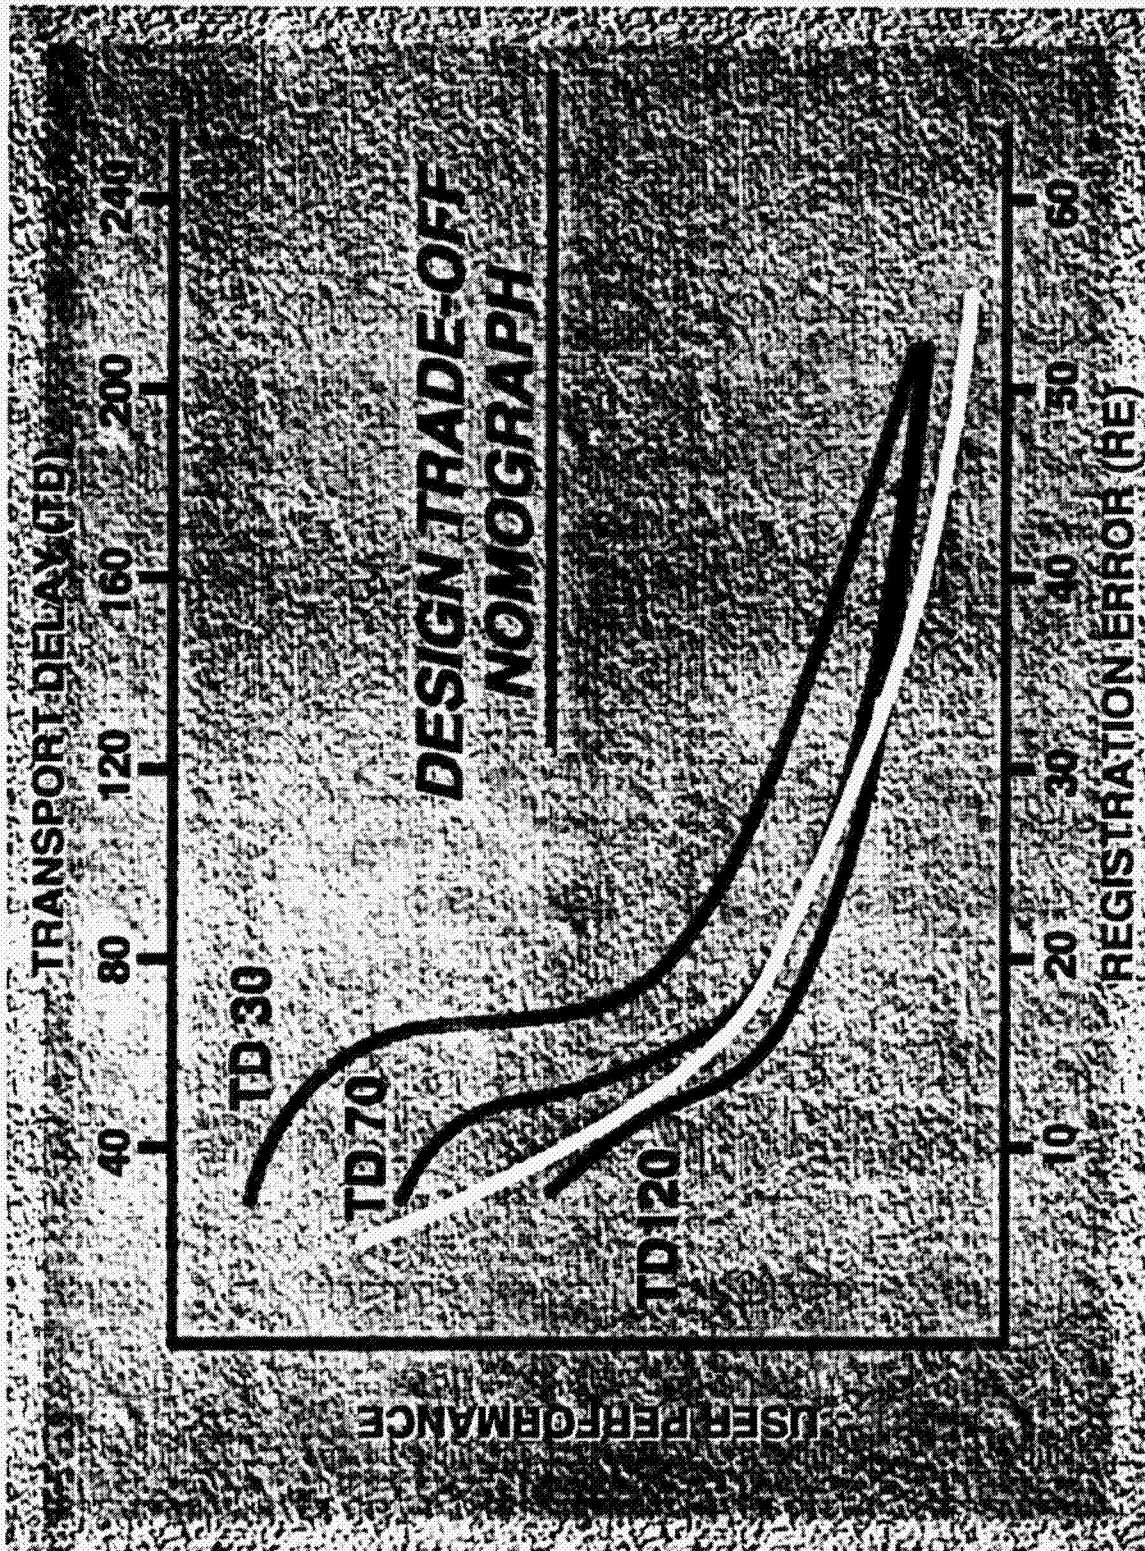


Figure 20

CLOSING REMARKS

As you can see, Armstrong Laboratory is engaged in a broad range of R&D research that involves both the development and exploitation of VR technology. We have a long history with the technology which has allowed us to achieve a high degree of understanding of the technical issues and challenges that need to be solved to insure its effective use in different application domains. We are very sensitive to the fact that VR technology by itself is neither desirable or undesirable. Rather, its value depends on how well we can unite it with the human user to produce successful products.



DETERMINING REQUIREMENTS AND CAPABILITIES OF VIRTUAL INTERFACES

Figure 21

**NEXT
DOCUMENT**

47

THE CAVE™ AUTOMATIC VIRTUAL ENVIRONMENT: CHARACTERISTICS AND APPLICATIONS¹

Robert V. Kenyon
Electronic Visualization Lab
University of Illinois at Chicago
Department of Electrical Engineering and Computer Science
Chicago, IL

INTRODUCTION

Virtual reality may best be defined as the wide-field presentation of computer-generated, multi-sensory information that tracks a user in real time. In addition to the more well-known modes of virtual reality – head-mounted displays and boom-mounted displays – the Electronic Visualization Laboratory at the University of Illinois at Chicago recently introduced a third mode: a room constructed from large screens on which the graphics are projected on to three walls and the floor.

The CAVE is a multi-person, room-sized, high-resolution, 3D video and audio environment. Graphics are rear projected in stereo onto three walls and the floor, and viewed with stereo glasses (Ref. 1). As a viewer wearing a location sensor moves within its display boundaries, the correct perspective and stereo projections of the environment are updated, and the image moves with and surrounds the viewer. The other viewers in the CAVE are like passengers in a bus, along for the ride!

“CAVE,” the name selected for the virtual reality theater, is both a recursive acronym (Cave Automatic Virtual Environment) and a reference to “The Simile of the Cave” found in Plato’s “Republic,” in which the philosopher explores the ideas of perception, reality, and illusion. Plato used the analogy of a person facing the back of a cave, live with shadows that are his/her only basis for ideas of what real objects are.

Rather than having evolved from video games or flight simulation, the CAVE has its motivation rooted in scientific visualization and the SIGGRAPH 92 Showcase effort. The CAVE was designed to be a useful tool for scientific visualization. The Showcase event was an experiment; the Showcase chair, James E. George, and the Showcase committee advocated an environment for computational scientists to interactively present their research at a major professional conference in a one-to-many format on high-end workstations attached to large projection screens. The CAVE was developed as a “virtual reality theater” with scientific content and projection that met the criteria of Showcase.

[™] CAVE is a registered trademark of the Regents of the University of Illinois.

¹ This research was supported by NSF Grant number IRI-9213822

GENERAL CHARACTERISTICS

The CAVE is a theater 10x10x9', made up of three rear-projection screens for walls and a down-projection screen for the floor (Figure 1). Electrohome Marquee 8000 projectors throw full-color workstation fields (1280x492 stereo) at 120 Hz onto the screens, giving between 2,000 and 4,000 linear pixel resolution to the surrounding composite image. Computer-controlled audio provides a sonification capability to multiple speakers. A user's head and hand are tracked with Ascension tethered electromagnetic sensors. Stereographics' LCD stereo shutter glasses are used to separate the alternate fields going to the eyes. A Silicon Graphics Onyx with three Reality Engines is used to create the imagery that is projected onto three or the four walls. The CAVE's theater area sits in a 30x20x13' light-tight room, provided that the projectors' optics are folded by mirrors (Ref. 2).

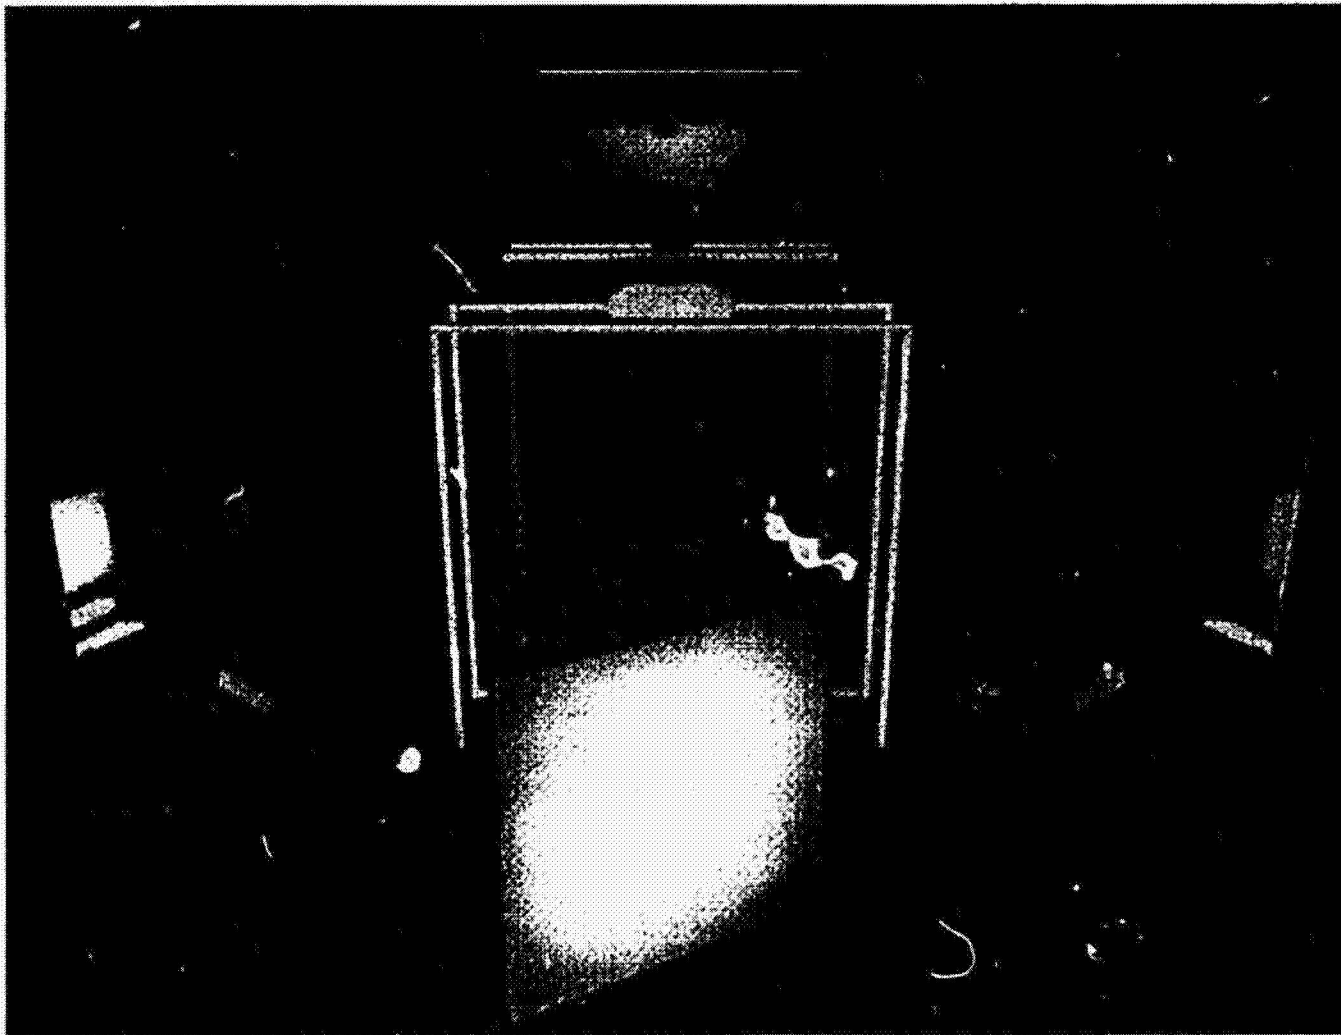


Figure 1. The CAVE

CAVE CHARACTERISTICS

Over the past two years, as we considered building the CAVE, there were several inherent problems with head-mounted virtual-reality technology to which we gave a great deal of thought:

- a. Simplistic real-time walk-around imagery
- b. Unacceptable resolution (the popular head-mounted displays offer resolution that is twice as bad as being legally blind)
- c. Difficulty of sharing experiences between two or more people
- d. Primitive color and lighting models
- e. No capability for successive refinement of images
- f. Too sensitive to rapid head movement
- g. No easy integration with real control devices
- h. Disorientation a common problem
- i. Poor multi-sensory integration, including sound and touch

The CAVE has the current capabilities and engineering characteristics:

- **Multi-person Virtual Environment**
- **Back projection onto 10'x10'x9' room**
- **Wide field-of-view**
- **High resolution color images**
- **Inside-out surround 3-D video presentation**
- **Off axis stereo projection**
- **Head and hand-tracked user interaction**
- **Co-existing real and virtual objects**
- **3-D Audio**
- **SGI Onyx with 3 Reality Engine 2s**
 - ⇒ **2 adjoining walls and the floor**
- **No force or tactile feedback (yet)**
- **Expensive**

Figure 2

VIDEO SYSTEM

The CAVE has an inside-out viewing paradigm where the design is such that the observer is inside looking out as opposed to the outside looking in (Ref. 3). The CAVE uses "window" projection where the projection plane and the center of projection relative to the plane are specified for each eye, thus creating off-axis perspective projection (Ref. 4). The correct perspective and stereo projections are based on cues returned by the Ascension position sensor attached to the Stereographics Crystal Eyes stereo shutter glasses. Each screen updates at 96 Hz or 120 Hz with a resolution of 1025x768 or 1280x492 pixels per screen, respectively. Two off-axis stereo projections are displayed on each wall. To give the illusion of 3-D, the viewer wears stereo shutter glasses that enable a different image to be displayed to each eye by synchronizing the rate of alternating shutter openings to the screen update rate. When generating a stereo image, the screen update rate is effectively cut in half due to the necessity of displaying two images for a 3-D image. Thus, with a 96 Hz screen update rate, the total image has a maximum screen update rate of 48 Hz. The CAVE has a panoramic view that varies from 90° to greater than 180° depending upon the distance of the viewer from the projection screens. The direct viewing field of view is about 100° and is a function of the frame design for the stereo glasses.

However, the reduction in resolution and update rate could be overcome with some design changes to the CAVE's current display system (Figure 3 bullet 1) or to future projector systems (Figure 3 bullet 2). For example, doubling the number of projectors per screen along with the number of graphics processors would restore the display to the original resolution (1024 horizontal lines) and update rate (96 Hz). To restore stereovision without shutter glasses the user would wear passive crossed polarizers with matching polarizers on the corresponding projector.

- **Electrohome high resolution projectors**
 - ⇒ **1024 x 1280 resolution**
 - ⇒ **Quick response time**
 - ⇒ **Fast green phosphor**
- **Future use of LCD lightvalves**
 - ⇒ **GII Wall at SuperComputing 95**
 - ⇒ **Crossed polarizers; not sequential video**
 - ⇒ **Polarizers in CAVE:**
 - **Folded optics problem**
 - **2 projectors per wall**
 - **Alignment**

Figure 3

VISUAL CHARACTERISTICS

Current VE applications are in some ways more ambitious and run on systems that have less computational power than current flight simulation applications. A chief attribute provided by most VE applications that can impact system performance, is user interaction with proximal virtual objects. To work effectively with objects at close range a user requires that the VE provide stereovision. This one necessity alone creates a series of constraints affecting the virtual environment. Stereovision requires that the user's current head position and orientation in the space be used so that the correct perspective views for each eye are generated. Without such information the 3-D world appears distorted. Consequently, the need to know head location forces the use of head tracking equipment that can compromise overall system performance in areas such as image update rate and lag (Ref. 5).

- **1024 x 768 Stereo resolution/screen**
- **2000 x 2000 Linear pixel resolution**
- **90° to 180° Horizontal Field-of-view**
- **100° Vertical Field-of-view**
- **Virtual object float inside CAVE**
 - ⇒ **Users can walk around objects**
- **Low screen brightness**
- **Accommodative stimulus at the screen**
- **Convergence stimulus on object**

Figure 4

AUDIO SYSTEM

At this time only directional sound is produced by the CAVE audio system but future plans call for 3-D audio production using Head-Related Transfer Function (HRTF). A MIDI synthesizer is connected via Internet/PC so, for example, sounds may be generated to alert the user or convey information in the frequency domain. Since the introduction of new systems to make the measurement of individual HRTFs more tractable, 3-D audio will soon be applied to the CAVE. However, at this time, only one person can be tracked and therefore the 3-D sound can only be correct for that person. This is a significant problem for systems that accommodate multiple users of the same environment such as the CAVE (Ref. 6).

- **6 Speaker system with controller**
 - ⇒ **General directional sound**
- **3-D Audio**
 - ⇒ **HRTF computed audio through earphones**
 - ⇒ **Head tracked for one person only**
 - ⇒ **Difficult for multiple users**

Figure 5

MAGNETIC TRACKING SYSTEM

Head and hand position are measured with the Ascension Flock of Birds six degree-of-freedom electromagnetic tracker operating at a 60 Hz sampling frequency for a dual sensor configuration. The normal and the augmented operation of the tracker in the CAVE is outlined in Figure 6. The transmitter is located above the CAVE in the center and has a useful operating range of 6 feet. Head position is used to locate the eyes to perform the correct stereo calculations for the observer. The CAVE's second position sensor is used to allow the viewer to interact with the virtual environment. Since this system is nonlinear and such nonlinearities can significantly compromise the virtual experience of immersion for the user, a calibration of the tracker system is needed. Nonlinearities caused by the metallic objects and electromagnetic fields created by other devices resident in and about the CAVE are compensated to within 1.5% by linearizing values returned by the head tracking system using a correction table containing calibrated positions in the CAVE (Ref. 7).

- **Dual Sensor 6 DOF magnetic tracking system**
- **Normal Range of operation 6 ft. radius**
- **Linear range reduced by distortions**
 - ⇒ **A 3 ft radius**
 - ⇒ **EM devices about the CAVE**
- **Extended linear range with calibration**
 - ⇒ **Nonlinearities reduced**
 - ⇒ **Accuracy improved for position**
 - ⇒ **Rotation calibration in progress**

Figure 6

CAVE CALIBRATION

The goal of the calibration procedure (outlined in Figure 7) is to correct for static position errors in the magnetic tracker. Metal structures near the tracker distort the magnetic field, so the CAVE screen frame is made of austenetic stainless steel which is non-magnetic and has a low conductivity. However, other components needed for the CAVE to function such as projectors and mirrors significantly distort the field. These distortions produce errors in the position component of the 6-degree-of-freedom magnetic tracker. By comparing the output with a custom-built ultra-sonic measuring system to the position reported by the magnetic tracker, a look-up table is created from the collected difference data and is used to interpolate for corrected values. The error of the resulting corrected magnetic tracker position is measured to be less than 5% over the calibrated range

- **CAVE filled with array of cubic calibration locations**
- **Magnetic/Sonic Calibration probe**
- **Probe placed inside cube**
- **X,Y,Z location recorded by sonic system**
- **Sonic data correlated to simultaneous Magnetic data**
- **Calibration array generated and data interpolated from lookup table**

Figure 7

CALIBRATION METHOD

The CAVE is first filled by a 3D stereo graphic image of 1-inch boxes on 1-foot intervals (Figure 8). A 1-inch cursor shows the position of the magnetic sensor which is placed atop the ultrasonic measurement device (UMD). A person wearing 3D glasses holds the UMD reasonably straight and moves it until the displayed cursor is inside of each box. The program records the position given by the magnetic sensor and the Onyx sends a signal to the PC to get the position measured by the UMD. This procedure continues until all the boxes in the tracker range inside the CAVE are thus sampled. In practice less than 400 points are collected, essentially all points in the center of the CAVE.

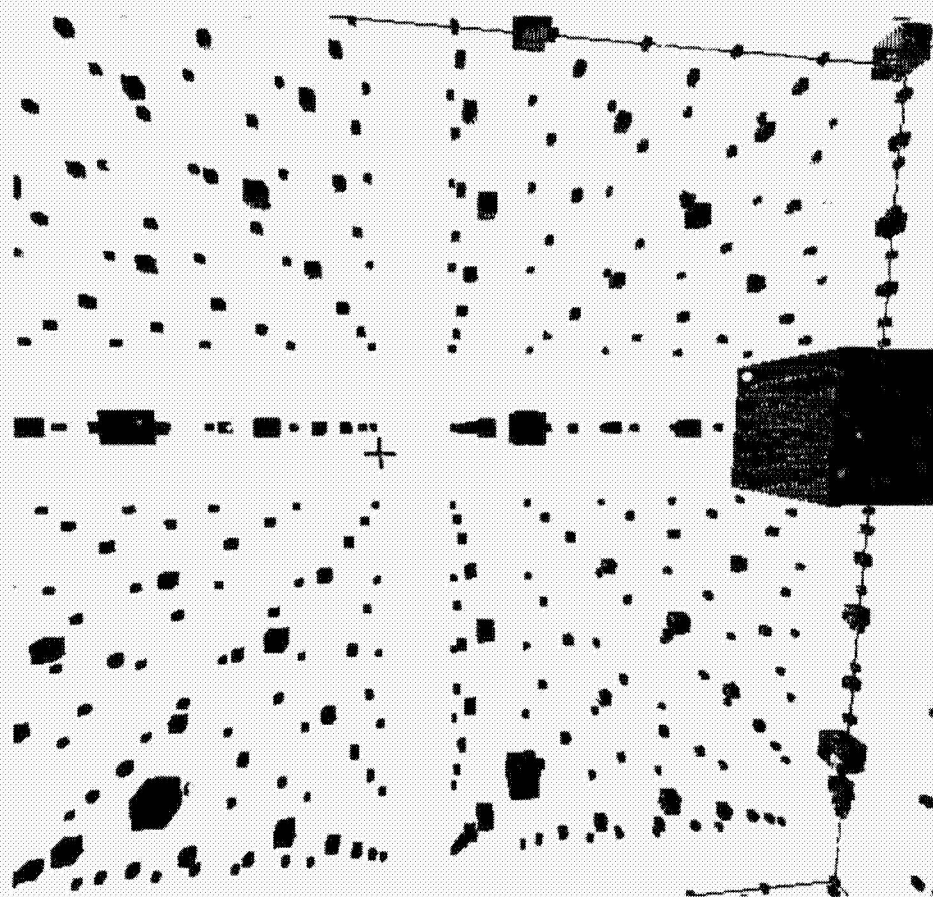


Figure 8

EFFECTS OF CALIBRATION

To measure residual errors after calibration we collected data at one foot intervals on half-foot centers instead of one foot interval on one foot centers. Therefore we measured residual errors half way between the calibration points. These results are shown in the figures below. They show a dramatic reduction in position error after the calibration is performed. These measurements of course depend on the accuracy of the UMD (less than 1.5% over 10 ft.). The maximum error before calibration is seen to be 4 ft. over a 10 ft. range (40%) (Figure 10). The error after calibrating is 0.27 ft. in the same 10 ft. range (2.7%). Similarly, the maximum error before calibration is 0.6 ft. in a 3 ft. range (20%). The error after calibrating is 0.13 ft. over the same range (4.3%). Clearly, this procedure is better at correcting larger errors than smaller ones, why this is true is not well understood at this point. Minimizing tracker latency is desirable in VR systems, so it is important that the correction computation does not substantially increase existing tracker latency. This linear interpolation method needs 30 additions and 72 multiplications for each correction. On the CAVE Onyx R4400 processor, the above calculation takes less than 10 microseconds. Since the theoretical minimum tracker latency is 21 milliseconds, adding 10 microseconds of delay is negligible.

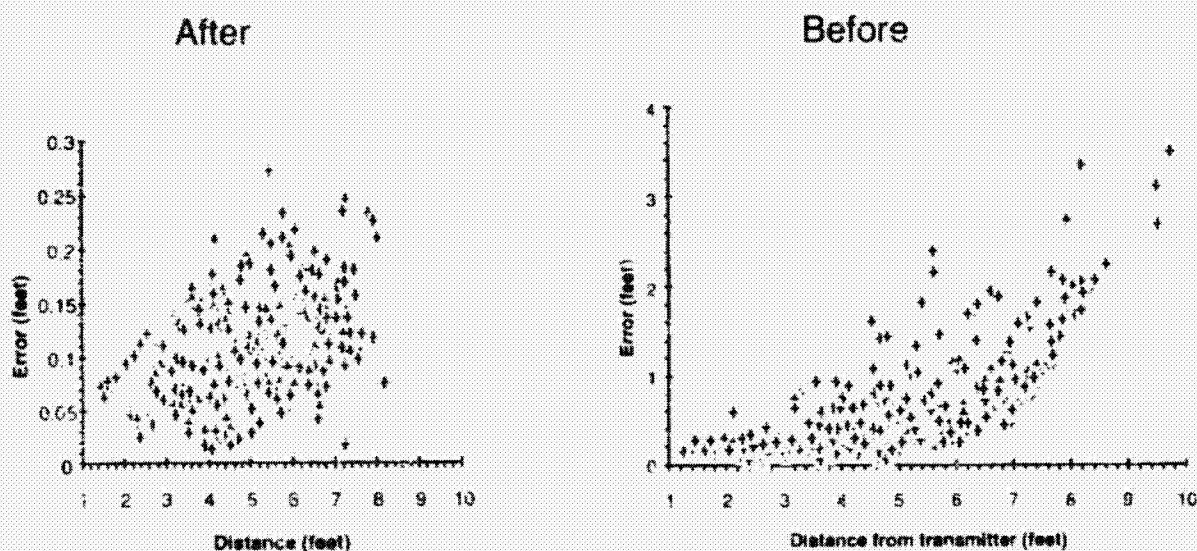


Figure 9

UPDATE RATES

As with most VR systems, the CAVE has a significant delay between the motion of the sensor and the resulting movement of the computer-generated scene. Currently, the minimum delay in the CAVE is in excess of 100 ms. Preliminary figures given of the characteristics in Figure 10 are:

1. Tracker to computer via RS232 line	=> 16 ms (Dual sensor Ascension Tracking System)
2. Serial port to shared memory area	=> 50 ms (Assuming no dedicated Onyx processor)
3. Shared memory area to rendering process	=> 20 ms (CAVE Library delay)
4. Rendering of the scene (variable)	=> 21 ms (Minimum for stereo)
TOTAL	107 ms

- **Stereo Image Update: 96 Hz**
- **Dual Sensor Magnetic Tracker: 60 Hz**
- **UNIX serial line handler: Bottleneck**

Figure 10

SYSTEM LAG

Some of the sources of system lag are outlined in Figure 11. We hope to reduce tracker delay to below 10 ms with improved serial interface cards that use IEEE 488 rather than RS232 protocol. The second largest item of 50 ms delay (in Figure 10) will decrease when a single processor is completely dedicated to the task of acquiring data from the serial port for use by the CAVE library routines. Without a dedicated processor, the delay for this process jumps to 80 ms. The reasons for this large delay are not clear at this time and software improvements will be needed to make this a more reasonable value. Regarding the CAVE library, we feel that further optimization will cut the delay to below 5 ms. Rendering time is somewhat more fixed given the graphics hardware. The display latency is set, however, by our choice of frame rate at 48 Hz which translates to a fixed 21 ms delay.

- **Pipeline Delays**
 - ⇒ **Serial Communication**
 - ⇒ **UNIX delays**
 - ⇒ **Software delays**
- **Current lag > 100ms**
- **ScramNet connection:**
 - ⇒ **Dedicated optical fiber network**
 - ⇒ **PC to Onyx direct memory connection**
 - ⇒ **Lag expected to drop below 90 ms**

Figure 11

CAVE DIRECTIONS

A clear interest in using CAVEs for collaborative virtual prototyping over distance has been voiced by our industrial clients and researchers in national laboratories. Fortunately, many of the barriers to effective use of this elaborate form of computer-mediated teleconferencing are similar to difficulties already under examination when attempting to use remote supercomputers as simulation and database servers.

In addition, we are making the CAVE both smaller and larger (Figure 14). The ImmersaDesk™ will bring the CAVE to the size of a drafting table. The Global Information Infrastructure (GII) project will expand the size of the CAVE to that so an entire audience may experience the immersion of a virtual environment in a familiar theater format.

- **CAVE to CAVE Interaction**
 - ⇒ **Long distance cooperative VE**
 - ⇒ **Collaborative virtual prototyping**
- **ImmersaDesk**
 - ⇒ **Smaller**
 - ⇒ **Less Expensive**
- **GII/Wall**
 - ⇒ **Larger Screen Size**
 - ⇒ **Higher Resolution**
 - ⇒ **Audience Immersion**

Figure 12

™ ImmersaDesk is a trademark of the Regents of the University of Illinois.

CAVE-TO-CAVE ISSUES

Some of the issues that we will be exploring in this area are summarized in Figure 13. Once again, latency is the issue, and given that latency over distance is unavoidable, compensation techniques must be developed. We, at this time, know the questions regarding variable latency in the CAVE-to-CAVE or CAVE-supercomputer-CAVE model; we do not have quantitative answers. We can, however, construct actual test situations and build measurement and assessment subroutines into our libraries.

Some of the issues that we are beginning to explore are:

1. What is the peak transfer rate of polygons (triangular meshes, in practice) from shared memory of the Onyx to multiple screens, and how do various raster manager boards affect this rate?
2. What is the peak transfer rate from striped disks? How fast can we get data over local FDDI and HIPPI networks? How to use ATM networks with existing and planned hardware?
3. What happens when distance is introduced, say from Chicago to Urbana? The distance, 140 miles, is less than 1 ms away at the speed of light. Where are the delays in the existing DS3-based networking, and will the vBNS OC3/12 improve the response time as expected?
4. How can we compress the data further than triangular meshes (one triple per polygon)?
5. How much synchronization is necessary between CAVE's to provide a useful collaborative virtual prototyping capacity? What is the tradeoff between maintaining local data at each CAVE site and sharing data in real time? Does this change if a supercomputer is providing the data?

- **Minimum necessary local compute power**
- **I/O system limitations**
- **Minimum network performance**
- **Latency vs. network state**
- **Optimal data compression**
- **Synchronization between CAVE's**
- **Maintaining local data at each CAVE site**

Figure 13

IMMERSADESK

We are developing ways to make the CAVE both smaller and more affordable. The "ImmersaDesk" is a drafting-table format virtual prototyping device (characteristics summarized in Figure 14). Using stereo glasses and sonic head and hand tracking, this projection-based system offers a type of virtual reality that is semi-immersive. Rather than surrounding the user with graphics and blocking out the real world, the ImmersaDesk features a 4x5' rear-projected screen at a 45° angle. The size and position of the screen give a sufficiently wide-angle view and the ability to look down as well as forward. The resolution is 1024 x 768 at 96Hz. It will also work in non-stereo mode at 1280 x 1024 at 60Hz.

We have learned from working with the CAVE that immersion is critical to achieving a workable virtual reality/prototyping system, and that immersion is dependent on being able to look forward and down at the display in such a way that the edges of the screen are not seen, or at least not prominent. Head-tracked stereo is important for virtual reality as well, although this can be easily achieved with a high-end workstation, desktop monitor and active stereo glasses. The ImmersaDesk allows the necessary wide angle of view and, because of its screen angle, the capability to portray forward and down views on one screen. Many researchers have developed stereo, even head-tracked monitor and projection-based wall virtual reality systems, but these do not allow down views and typically have narrow angles of view.

The ImmersaDesk is a derivative of the CAVE system, being an excellent development environment for the CAVE, and also a stand-alone system. The ImmersaDesk prototype is 100% software compatible with the CAVE libraries and interfaces to software packages like Sense8's World Toolkit and SGI's Performer/Inventor, as well as visualization packages like AVS and IBM Data Explorer. Interfaces to industry standard CAD output files are also provided via these packages.

Since the CAVE/ImmersaDesk libraries have been much used to view high-bandwidth supercomputer output, the VTC/networking features will also permit the interactive shared steering of computations and the querying of databases by a number of people. This neatly combines the best of video communications with the best of simulation computing and the high-end of virtual reality interactive 3D visualization.

- **Drafting-table format**
- **Relatively small: 6' x 8' footprint**
- **Requires:**
 - ⇒ **Single pipe Onyx**
 - ⇒ **One projector**
 - ⇒ **No Architectural Modifications**
 - ⇒ **compatible with CAVE lib**
- **Interfaces to:**
 - ⇒ **Sense8's World Toolkit**
 - ⇒ **SGI Performer/Inventor**
 - ⇒ **CAD output files**

Figure 14

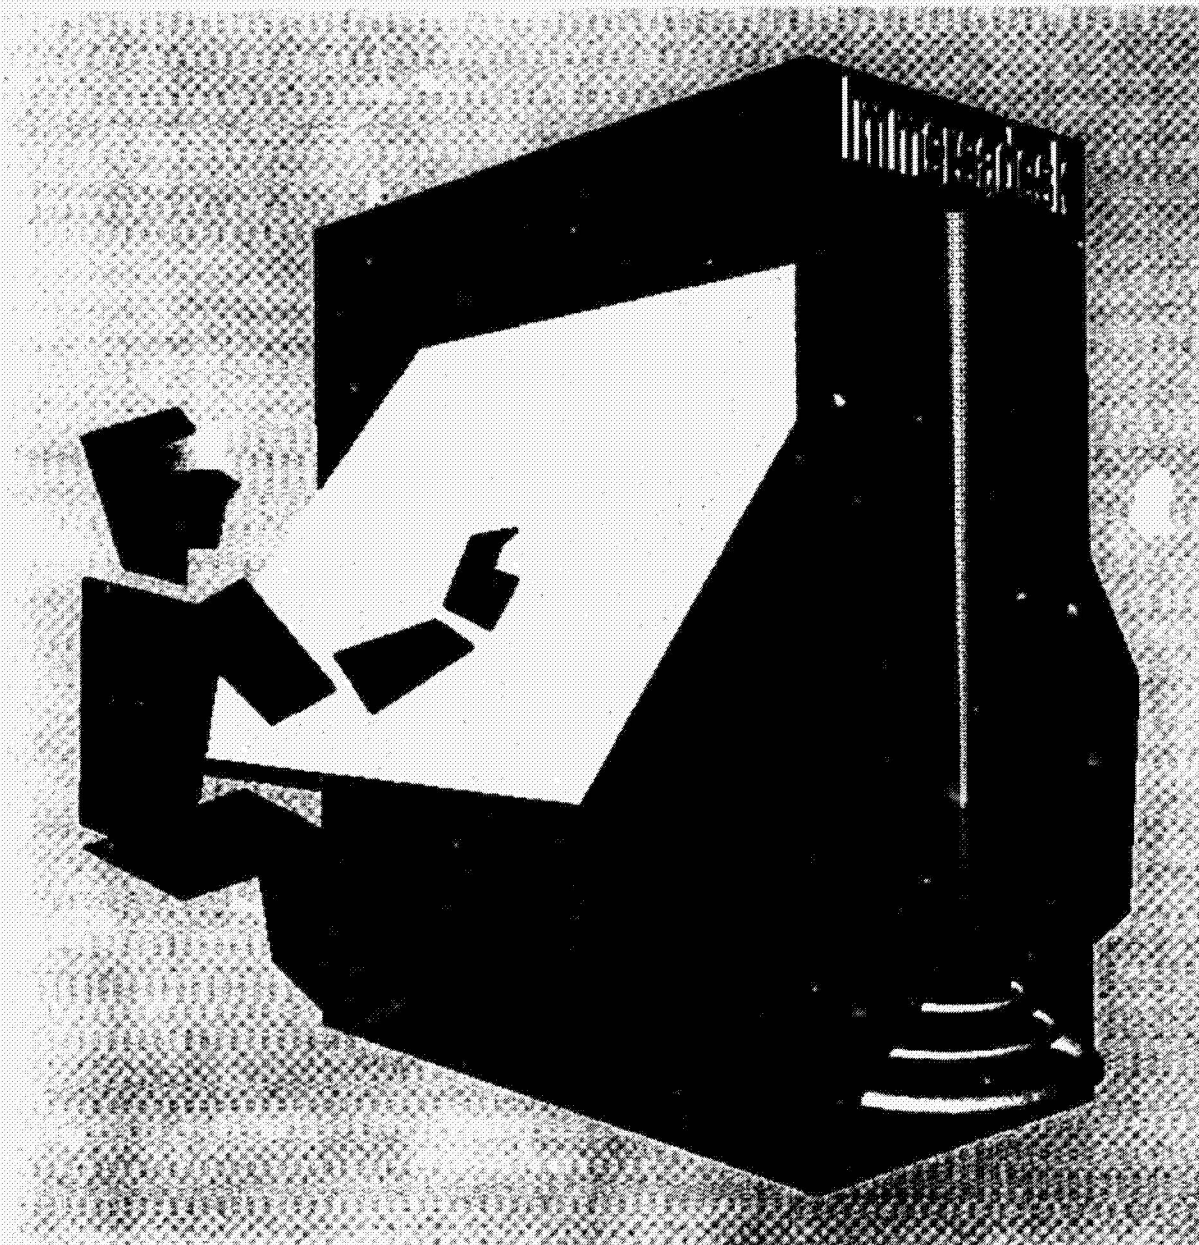


Figure 15. The ImmersaDesk. The user is seated in front of the slanted screen (left of the figure) and wearing stereo shutter glasses to obtain 3-D images. The projector and optics are contained within the box holding up the projection screen.

GII/WALL

The Super Computing'95 Global Information Infrastructure (GII) Testbed event provides a venue for inter-active 2D and 3D demonstrations of National Challenges and Grand Challenges – remotely computed in a scientist's numerical laboratory and then transmitted over high-speed networks for presentation in San Diego. The characteristics of this large format CAVE is outlined in Figure 15.

The GII/Wall is a large-screen, high-resolution (1600 x 2048) stereo projection display. A non-stereo PowerWall was developed by Paul Woodward's group at the Army High Performance Computer Research Center at University of Minnesota for the Silicon Graphics' booth at SC'94.

The GII/Wall uses four Reality Engines spread across two Power Onyxes to achieve high-resolution, high-intensity, passive-stereo images. One Onyx controls the top half of the screen and the other controls the bottom half (each at 1600 x 1024). The top and bottom are each driven by two Reality Engines displaying polarized projected images for each eye. Throwaway polarized glasses can be used by the audience instead of the active stereo glasses used in the CAVE and ImmersaDesk systems. The GII/Wall is much more suited for audiences, and although it uses a lot of computers and projectors, the number of people it reaches per unit time is far greater than either the CAVE or the ImmersaDesk.

The GII/Wall achieves its immersion by wide-screen projection, but does not allow, unfortunately, a way to look down, a problem with any normal audience seating arrangement. (Note that the angle of Omnimax/ Imax theater seating addresses this problem by steeply pitched seating). The developers are currently experimenting with large-area types of tracking, mindful of the fact that it is only possible to track one person at a time, not an audience. The GII/Wall is appropriate for applications in which high-resolution telepresence is the goal rather than audience participation.

- **Goal:**
 - ⇒ **High resolution telepresence for large audiences**
 - ⇒ **Not for audience participation**
- **Large Screen/Wall Viewing**
- **High Resolution (1600 x 2048)**
- **High Intensity (lightvalves)**
- **Passive Stereo**
- **2 Power Onyxes:**
 - ⇒ **Each with 2 SGI Reality Engines**
 - **One Onyx Each for top and bottom image**
- **Networked scaleable computing**

Figure 16

SUMMARY

In summary, an alternate form of a Virtual Environment presentation system has been described. The general characteristics of its visual, auditory, tracking, and image generation systems have been detailed. Specific problems associated with this system have been addressed and effective solutions have been shown.

In addition, two derivatives of this system have been presented: The ImmersaDesk(TM) and the Global Information Infrastructure Wall. The first represents a smaller, less expensive, multi-person immersive system. The second, a larger, audience style, presentation format of an immersive environment.

Finally, interconnection of all these VE systems was discussed with the goal of collaborative virtual prototyping for science and industry.

**NEXT
DOCUMENT**

59-62
62+8

VIRTUAL REALITY FOR AUTOMOTIVE DESIGN EVALUATION

**George G. Dodd, Head
Analytic Process Department
General Motors Research and Development Center
Warren, MI**

INTRODUCTION

A general description of Virtual Reality technology and possible applications was given from publicly available material. A video tape was shown demonstrating the use of multiple large-screen stereoscopic displays, configured in a 10' x 10' x 10' room, to allow a person to evaluate and interact with a vehicle which exists only as mathematical data, and is made only of light. The correct viewpoint of the vehicle is maintained by tracking special glasses worn by the subject. Interior illumination was changed by moving virtual light around by hand; interior colors are changed by pointing at a color on a color palette, then pointing at the desired surface to change. We concluded by discussing research need to move this technology forward.

BACKGROUND

Many of us have seen magazine covers, such as Automotive Industries, describing virtual reality. Today I want to point out to you that virtual reality is not virtual, it's real and being used right now.

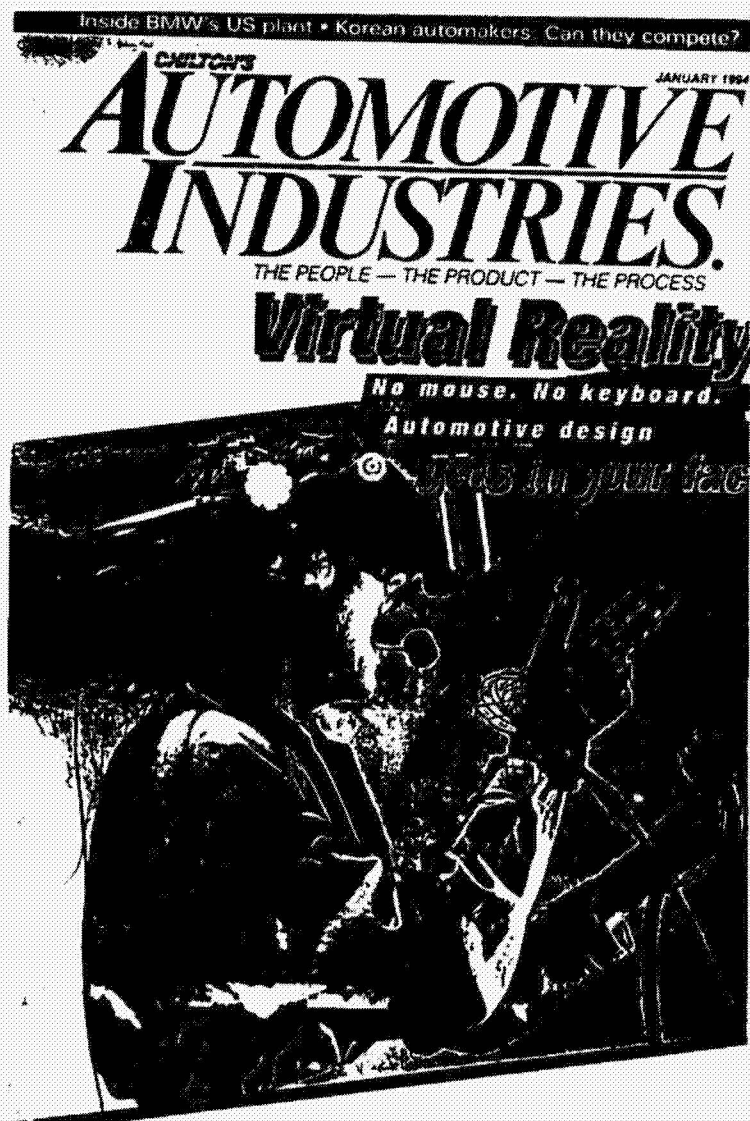


Figure 1.

AGENDA

In this presentation, we'll discuss virtual reality and how it can be used in vehicle design.

What is Virtual Reality?

How can it be used in design?

Figure 2.

VIRTUAL REALITY PROMISE

Virtual reality has a great deal of promise. If I hear something, I will probably forget it, because it goes in one ear and out the other; if I see it, I may remember; but if I do it, I will understand. Virtual reality allows us to experience an environment or experience a process with "us" as key players.

PROMISE OF VR TECHNOLOGY

If I hear . . . I will forget.

If I see . . . I may remember.

If I do . . . I will understand.

Old Chinese Proverb

Figure 3.

DEFINITIONS

What is virtual reality? Jason Lanier said that it is, "A computer generated, interactive, three dimensional environment which a person is immersed."

The virtual domain provides: (1) a real-time, viewer-centered system in which the viewer is a part of the display and the interaction; (2) a head-tracked perspective so that the image shifts as the viewer moves so that he accurately sees it from any point of view; (3) a large field of view, so the peripheral vision is encompassed in the field of view; (4) interactive control, so that you can control the images of the 3D stereo display.

And last, what is immersion? It is a state of being deeply engaged or immersed in the process. This is what virtual reality allows you to do.

What is Virtual Reality?

A computer-generated, interactive, three-dimensional environment in which a person is immersed

Jason Lanier, founder of VPL

What is a Virtual Domain?

A VR system that provides a real-time, viewer-centered, head-tracked perspective with a large field of view, interactive control, and 3-D stereo display

Carolina Cruz-Neira, EVL

What is Immersion?

The state of being deeply engaged

Webster

Figure 4.

1950'S STEREO VIEWING

Back in the 1950's, people would go to movies wearing glasses as shown in this picture. However, virtual reality is much more than watching a picture in 3D, it is experiencing and being involved in the picture. It surrounds you, with you as the center of attention.



Figure 5.

VIRTUAL REALITY PROTOTYPING

Virtual reality allows us to have a virtual prototype, which is a 3D computer model made completely with light with which designers and engineers can interact in the same ways they would interact with physical models, but at greatly reduced cost, at greatly increased speed of implementation, and with more flexibility and agility than you can find with physical models.

VIRTUAL PROTOTYPING

A virtual prototype is a 3-D computer model made completely with light that designers and engineers can interact with in some of the same ways they would a physical model, but . . .

- . . . at greatly reduced cost**
- . . . at greatly increased speed of implementation**
- . . . with more flexibility and agility**

Figure 6.

VIRTUAL REALITY KEY CONCEPTS

This slide illustrates the key concepts behind virtual reality. At the graphic terminal, a person designs something such as a seat, which becomes a number of mathematical equations and databases inside the computer. There are three ways to experience the results of the design: (1) the boom; (2) the head-mounted display; and (3) a CAVE.

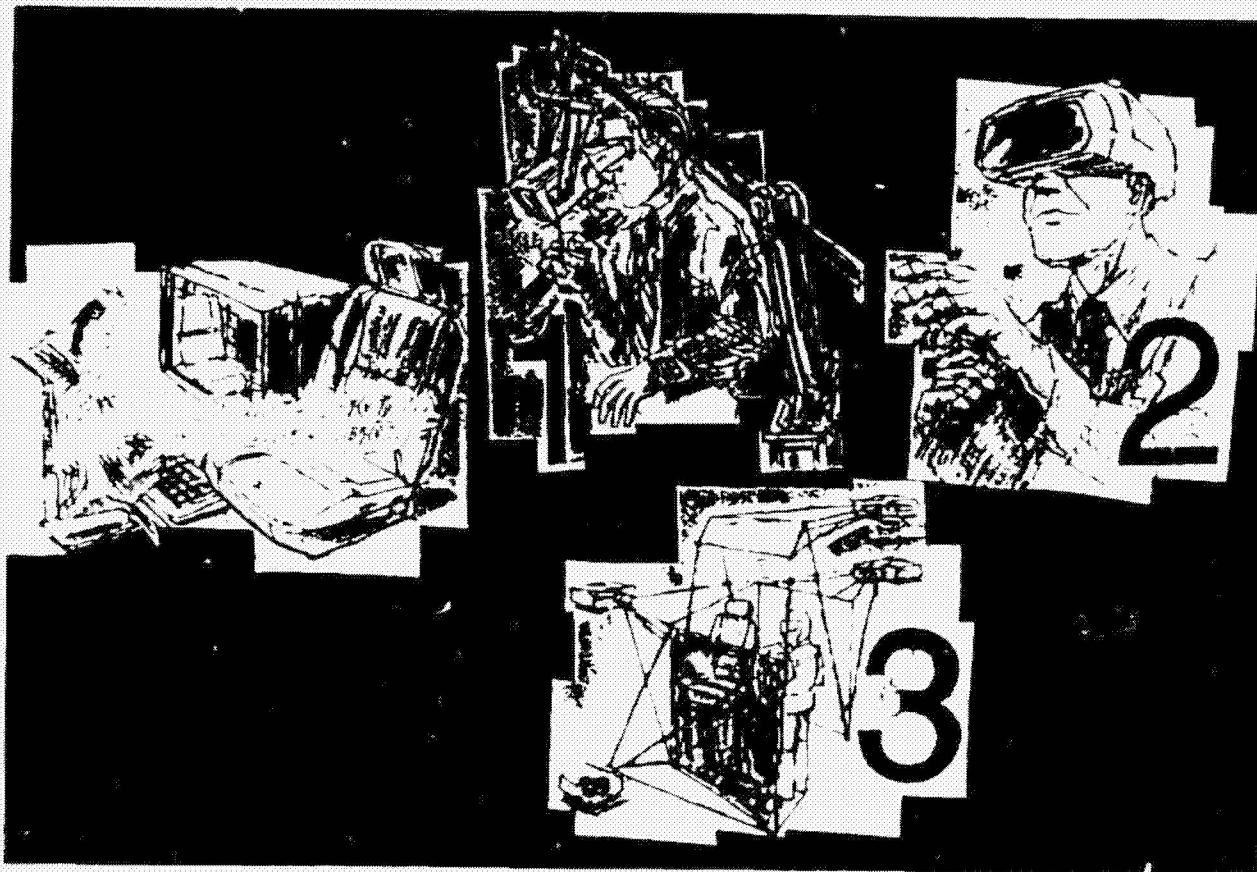


Figure 7.

BOOM DESCRIPTION

The boom is a device that you hold up to your eyes, like you would hold a pair of binoculars. You move around and the computer knows your position/orientation through a linkage mechanism connecting back to the computer. You change the distance to and from the object by pushing one or two buttons on the handles of the boom. You can control the interaction with the object through a data glove, which is then represented on the picture inside the viewing area.



Figure 8.

BOOM CHARACTERISTICS

The characteristics of a boom are as follows:

- It has a very narrow field of view, roughly the same as you would have by looking through a pair of binoculars
- It has accurate and fast tracking, but is a cumbersome periscope with manual sets of linkages to give you limited range of movement
- Virtual hand images are superimposed on the display
- However, the person using the boom is completely isolated from physical surroundings
- It is not very intuitive to use because of the short range and the push buttons
- You can have human factors problems such as dizziness, because you are not tied into the real world. If the image does not track as speedily as your head moves, the delay in movement can cause seasickness.

BOOM CHARACTERISTICS

- **Field of view narrow – 90 ... 120 degrees**
- **Accurate, fast tracking**
- **Cumbersome periscope, manual set of linkages**
- **Virtual hands image superimposed on display**
- **Completely isolated from physical surroundings**
- **Not intuitive to use**
- **Human factors (dizziness)**

Figure 10.

VIRTUAL REALITY HAND DESCRIPTION

This is a picture of a data glove being used in a surgical operation tied to a boom. Note the artificial rendition of the hand.

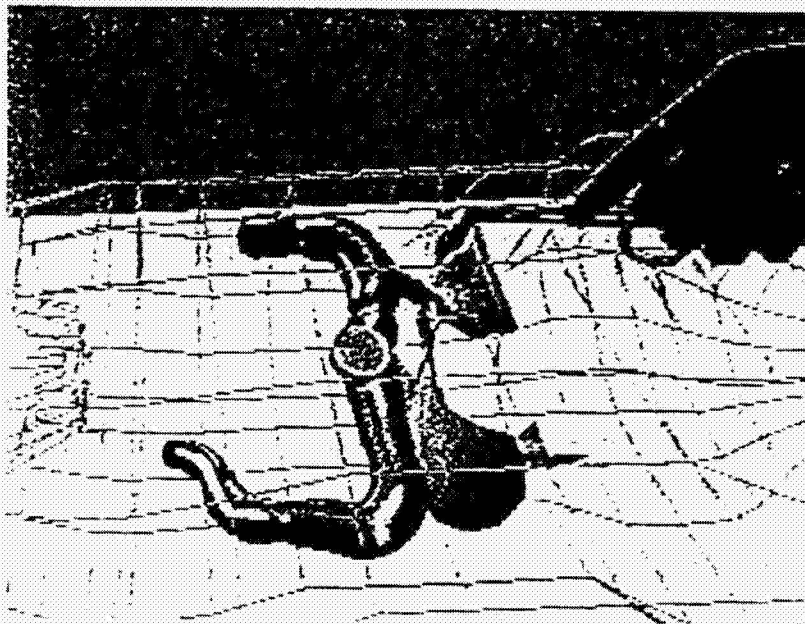


Figure 9.

HEAD-MOUNTED DISPLAY DESCRIPTION

The head-mounted display is a device that is mounted on the head, in which two TV images are projected. Weighing between five and seven pounds, it feels like a football helmet, but it is very cumbersome to wear. Because of the inertia of the weight, when you move your head around you sometimes have an inertia effect that doesn't feel very natural.

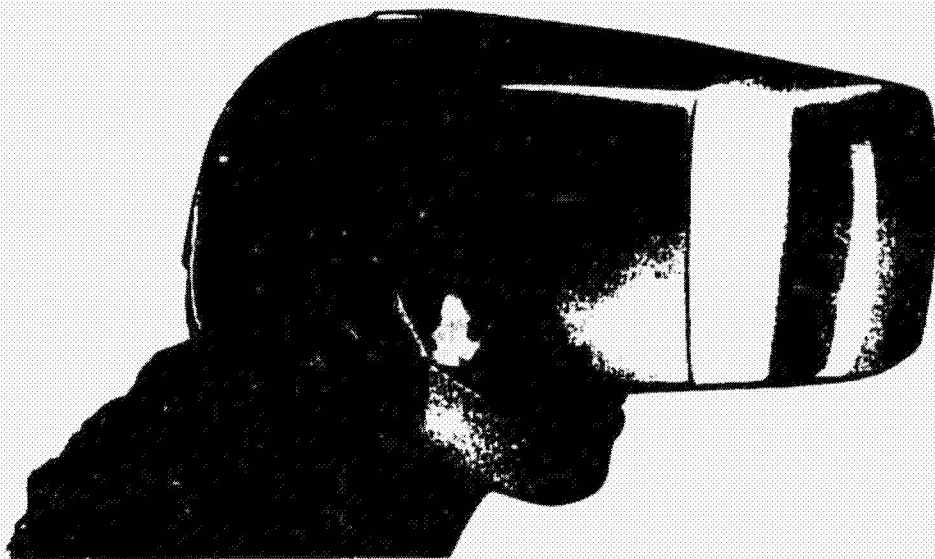


Figure 11.

HEAD-MOUNTED DISPLAY CHARACTERISTICS

The head-mounted display characteristics are shown here. In general, head-mounted displays have a wider field of view than the boom, and hands are free to interact. However, we have to display images of the hands on the display within the football helmet-like device. The helmet is cumbersome with tethered wires. Again, you're completely isolated from physical surroundings, with the accompanying human factors problems which could occur.

HEAD-MOUNTED DISPLAY CHARACTERISTICS

- **Field of view – 100 ... 140 degrees**
- **Hands free to interact**
- **Cumbersome helmet, tethered wires**
- **Completely isolated from physical surroundings**
- **Human factors (dizziness)**

Figure 12.

CAVE DESCRIPTION

The CAVE is a three-dimensional cubic room. At the General Motors Research & Development Center, it is a room 8' X 8' X 8'. On each of the three walls and the floor is a projected image. When you are inside the CAVE, you get a complete sense of immersion in that you are surrounded by the images made completely with light. Because of the projection onto the floor, the images can rise out of the floor. The person wears lightweight stereo glasses and the head tracking is controlled by a sensor which knows the position of the head within the room and can adjust the picture accordingly.

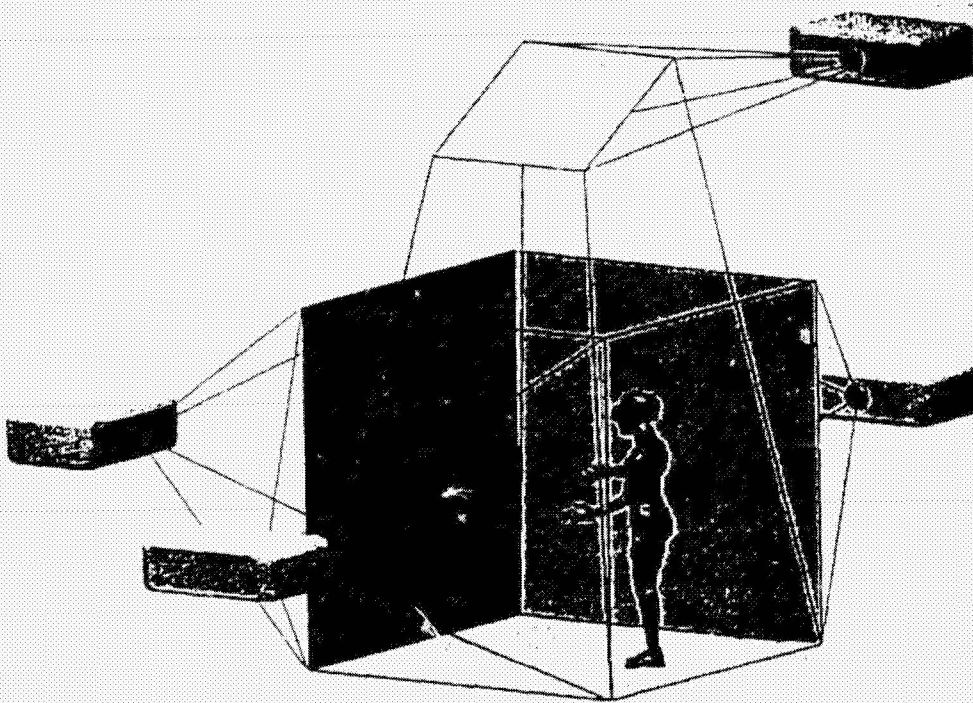


Figure 13.

CAVE VISUALIZATION

This picture shows one of our researchers, Randy Smith, with a pair of head-mounted display glasses, standing near a virtual image of a vehicle.

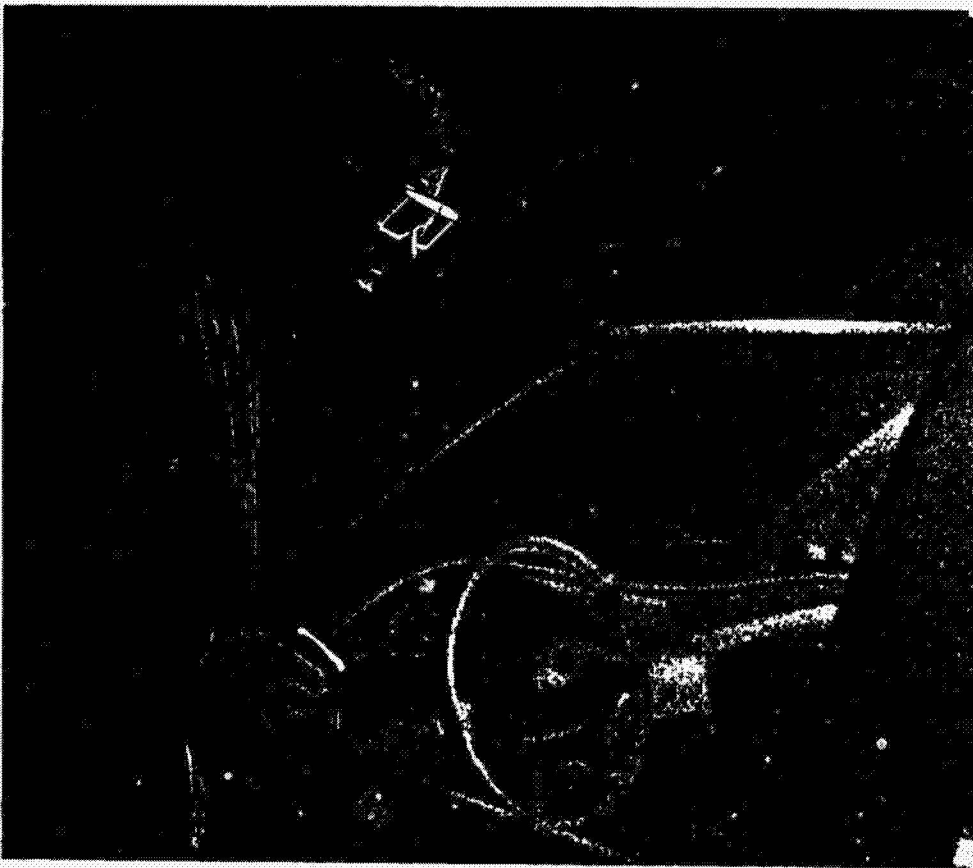


Figure 14.

CAVE CHARACTERISTICS

The CAVE gives you a complete full field of view, in which you are totally immersed. It is the least intrusive of all three virtual reality approaches in that you only wear eye glasses. You don't need to model and track the hands because you can see them as you are involved within the environment. One new characteristic of the CAVE that you don't find in the other two environments is that you can mix virtual and physical objects and do joint simulations of both types of medium. In addition, the CAVE allows for multiple participants to look at and interact with the display. Head tracking applies to one person only, but others standing close to him will see about the same image.

CAVE CHARACTERISTICS

- **Field of view full – Total immersion**
- **Least intrusive – requires wearing glasses only**
- **Don't need to model and track hands / fingers since you can see them**
- **Mix virtual and physical objects**
- **Multiple participants (perspective set for 1 only)**
- **Hands free to interact**

Figure 15.

HOW VIRTUAL REALITY IS USED IN DESIGN

Virtual reality can be used four ways in the design process: (1) to assess the overall impact of the design, (2) to conduct visual quality inspection, (3) to examine human factors and conduct human factors evaluation of the product, and (4) to study packaging.

How Can Virtual Reality Be Used in Design?

Overall impact of design

Visual quality inspection

Human factors evaluation

Packaging studies

Figure 16.

DESIGN IMPACT

Design impact is illustrated by modeling the interior of the vehicle. This enables flexible evaluation and modification of a design. By modeling the interior, you can then do "what if" studies visually. You can look for occlusion and accessibility of instruments. You can study the 'A' pillar, that's the pillar between the door and the windshield, for obscuration and exterior vision. You can study the effects of occupant height variation and overall interior styling of the vehicle.

DESIGN IMPACT

A computer model of an interior concept enables flexible evaluation and modification of a design

“What-if” studies can be performed to visualize:

- **Occlusion and accessibility of instruments**
- **‘A’ pillar obscuration and exterior vision**
- **Effects of occupant height variation**
- **Interior styling**

Figure 17.

VISUAL QUALITY INSPECTION

This image is made completely in a virtual world. The entire vehicle and the surrounding laboratory are all computer generated. No physical models were utilized in designing and displaying this picture.

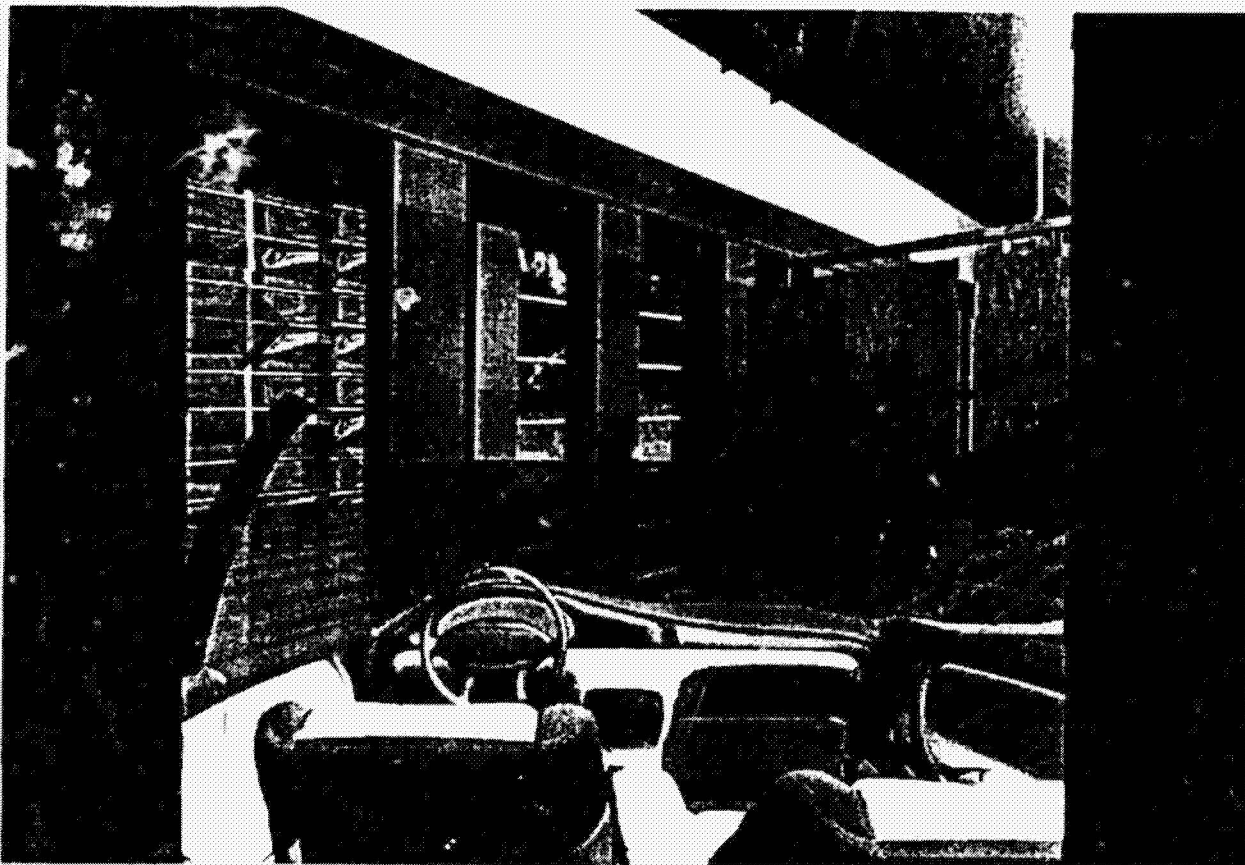


Figure 18.

HUMAN FACTORS EVALUATION

Human factors evaluation can be done by immersing the viewer in the display projected in the CAVE.

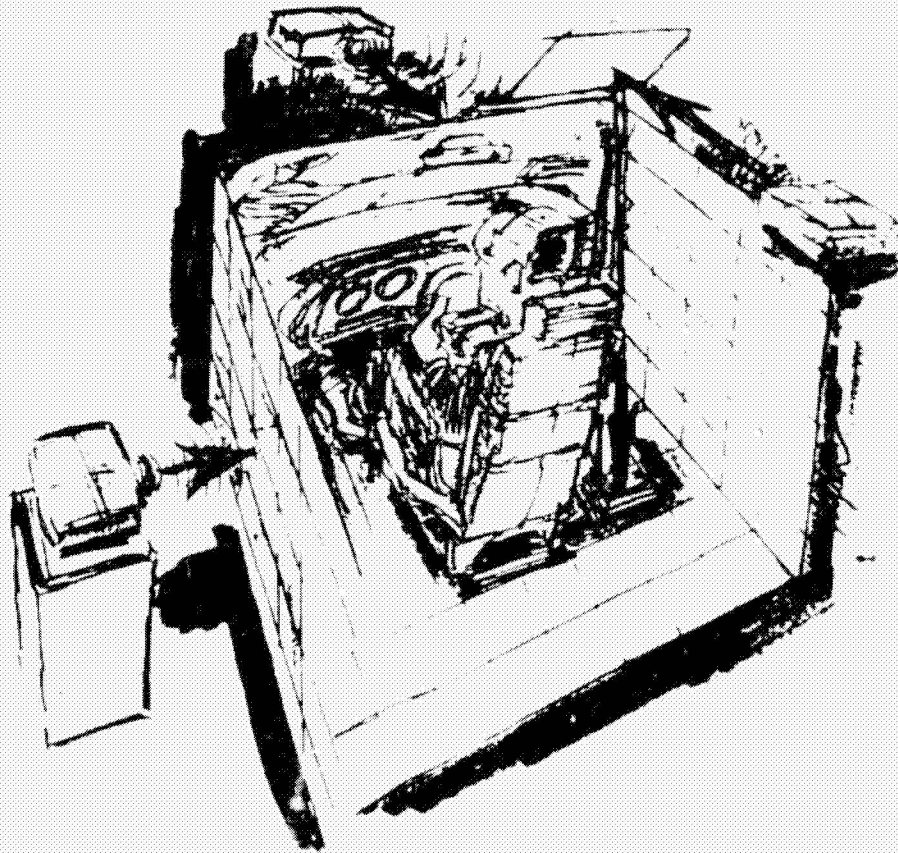


Figure 19.

PACKAGING STUDIES

Packaging studies can be conducted by assembling the parts of the vehicle, again designed within the computer CAD system without physical models, to study the packaging effects. You can fly through the object and look for interferences and visually inspect the assembly.

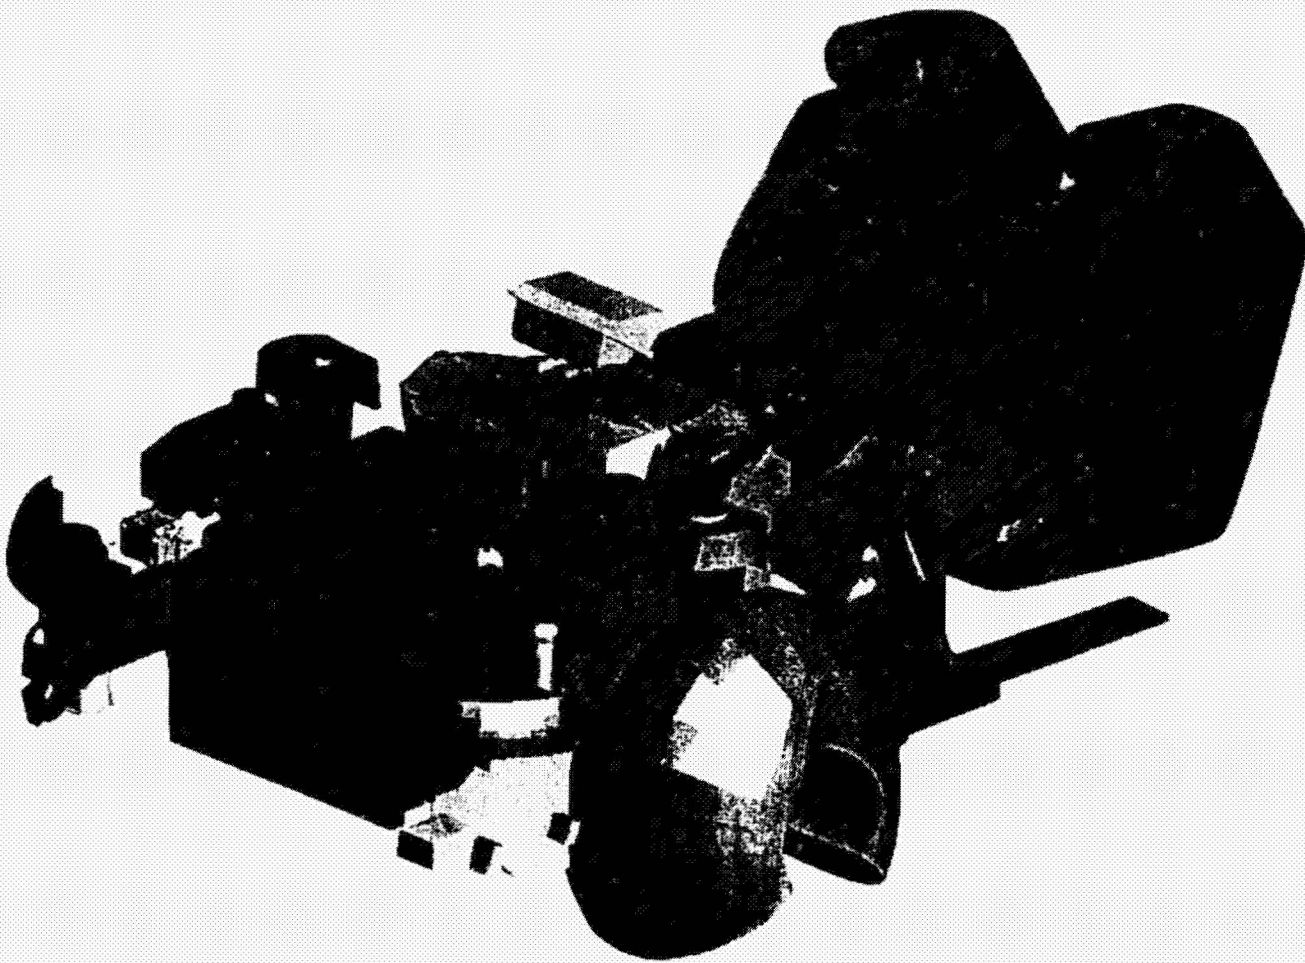


Figure 20.

OTHER APPLICATIONS

Virtual reality can also be used for other applications, such as scientific visualization from a dynamic viewpoint. In crash simulation, for example, you can sit inside a vehicle and view the crash from the inside, to determine how metal moves, how the steering wheel might move and make certain that the occupant would be safe in that kind of situation. You can move your viewpoint and study suspensions and mechanisms. In painting, you can study spray dispersal patterns to make certain that you get uniform spray on a piece of sheet metal, especially around curves and the edges of the metal. NASA Ames is using virtual reality for virtual wind tunnels to study aerodynamic effects.

OTHER APPLICATIONS

■ Scientific visualization – Dynamic viewpoint

- **Crash simulation – Full-size, viewed from inside**
- **Suspensions and mechanisms**
- **Metal forming**
- **Spray dispersal patterns**
- **Virtual wind tunnel (NASA Ames)**

Figure 21.

KEY FEATURES OF VISUALEYES

GM R&D's VisualEyes system has four key features: It has a human-in-the-loop simulation, in which the human is immersed in the scene. To do this, we are using the CAVE approach with head tracking to correct their perspective view. In fact it is so good, it allows you to stand up and walk around the interior of the vehicle, to look outside over the vehicle, and so forth. We can show models from any math data which meet GM's C4 standards, which are CGS, Unigraphics, Alias, and IGES.

GM R&D's VisualEyes – KEY FEATURES

- **Human-in-the-loop simulation**
- **CAVE approach**
- **Head tracking**
 - **Correct perspective view, “walk-around”**
- **Shows models from any math data meeting GM's C4 standards**
 - **CGS, UG, Alias, IGES**

Figure 22.

REQUIRED SOFTWARE AT RESEARCH & DEVELOPMENT

There are several R&D areas in which virtual reality needs more help. We need better software for faster rendering of complex scenes with lots of detail and better control over the detail. People want to have images that track with the head, not slightly behind, which is still the case in our very high resolution images. In addition we need to consider new human interface paradigms for 3D design. How could you sit inside a CAVE and design while you're inside the CAVE is one problem. How might you paint on the wall and have that relate to a display of the design on that wall is a related interesting problem.

R&D Areas

SOFTWARE

- **Faster rendering with lots of data; detail control**
- **New human interface paradigms for 3D design**

Figure 23.

REQUIRED HARDWARE AT RESEARCH & DEVELOPMENT

Looking at hardware R&D areas, we need to have tracking which is more precise, faster, multi-point, and with no wires. This permits you to move your head around, turn, bend up and down and have absolute realistic tracking of the eye point. We need high resolution, wide field displays and stereo immersive displays. It would be nice, for example, to use high definition television techniques to get much better quality display than we have today. And last, we need to consider better hardware for tactile input and response so we can simulate humans touching and feeling parts of the display.

R&D Areas

HARDWARE

- **Tracking ... more precise, faster, multi-point, no wires**
- **High-resolution, wide field-of-view stereo immersive displays**
- **Tactile input and response**

Figure 24.

**NEXT
DOCUMENT**

A Synthetic Design Environment for Ship Design

Richard R. Chipman
Science Applications International Corporation
McLean, VA

REQUIRED SYSTEMS AT RESEARCH & DEVELOPMENT

Looking at the system area, calibration techniques for accurate engineering and design work needs to be improved. Right now it takes a considerable amount of time to calibrate displays for realistic human simulation. As we have said before, we need better human-in-the-loop, real-time performance so that we can absolutely simulate the immersive effect in real time. On a more biological note, we need human perceptual studies conducted to determine the limits and requirements of how the psychology of the human eyes and ears interact to understand and interact with realistic displays.

R&D Areas

SYSTEM

- **Calibration for accurate engineering and design work**
- **Human-in-the-loop real-time performance**
- **Human perceptual studies to determine limits, requirements**

Figure 25.

0-32
6249

A SYNTHETIC DESIGN ENVIRONMENT FOR SHIP DESIGN†

~~1-29-82~~
12-1-82

Richard R. Chipman *
Science Applications International Corporation
McLean, Virginia

INTRODUCTION

Rapid advances in computer science and information system technology have made possible the creation of synthetic design environments (SDE) which use virtual prototypes to increase the efficiency and agility of the design process. This next generation of computer based design tools will rely heavily on simulation and advanced visualization techniques to enable integrated product and process teams to concurrently conceptualize, design and test a product and its fabrication processes. This paper summarizes a successful demonstration of the feasibility of using a simulation based design environment in the shipbuilding industry.

As computer science and information science technologies have evolved, there have been many attempts to apply and integrate the new capabilities into systems for the improvement of the process of design. These systems go by names like Computer-Aided Design (CAD), Computer-Aided Manufacturing (CAM), Computer-Aided Engineering (CAE), and Computer-Aided Systems Engineering (CASE). We see the benefits of those efforts in the abundance of highly reliable, technologically complex products and services in the modern marketplace. Furthermore, the computer-based technologies have been so cost effective that the improvements embodied in modern products have been accompanied by lowered costs.

Today the state-of-the-art in computerized design has advanced so dramatically that the focus is no longer on merely improving design methodology; rather the goal is to revolutionize the entire process by which complex products are conceived, designed, fabricated, tested, deployed, operated, maintained, refurbished and eventually decommissioned. By concurrently addressing all life-cycle issues, the basic decision making process within an enterprise will be improved dramatically, leading to new levels of quality, innovation, efficiency and customer responsiveness. By integrating functions and people within an enterprise, such systems will change the fundamental way American industries are organized, creating companies that are more competitive, creative and productive.

†This work was sponsored by the Advanced Research Projects Agency, Maritime Systems Technology Office.

* Project Manager, Industrial Technologies Division

TRADITIONAL PRODUCT DESIGN CYCLE

As indicated in Fig. 1 below, the product life cycle consists of five phases: (1) requirements definition, (2) concept development and design, (3) manufacturing, (4) test and verification, and (5) operation. (The operational phase can be thought to include product utilization, maintenance, refurbishment and disposal.) Although the traditional, narrow view of product development focused on phases (2) through (4), today it is well recognized that the product development cycle must include both the front-end, wherein some basic trades between user needs and system capabilities are made, and the tail-end, which contains much of the total life cycle cost and all of the value delivered to the customer.

Traditionally, product development has been a sequential process, starting with a perceived customer need that is translated into requirements and ending with a product in operation, hopefully satisfying the perceived need. As shown in the figure, iterations of the design can take place at any point along the process; but changes made to the product late in this cycle are inevitably costly. One of the basic objectives of SDE is to avoid such costly redesign by enabling designers to concurrently address all phases of the life cycle.

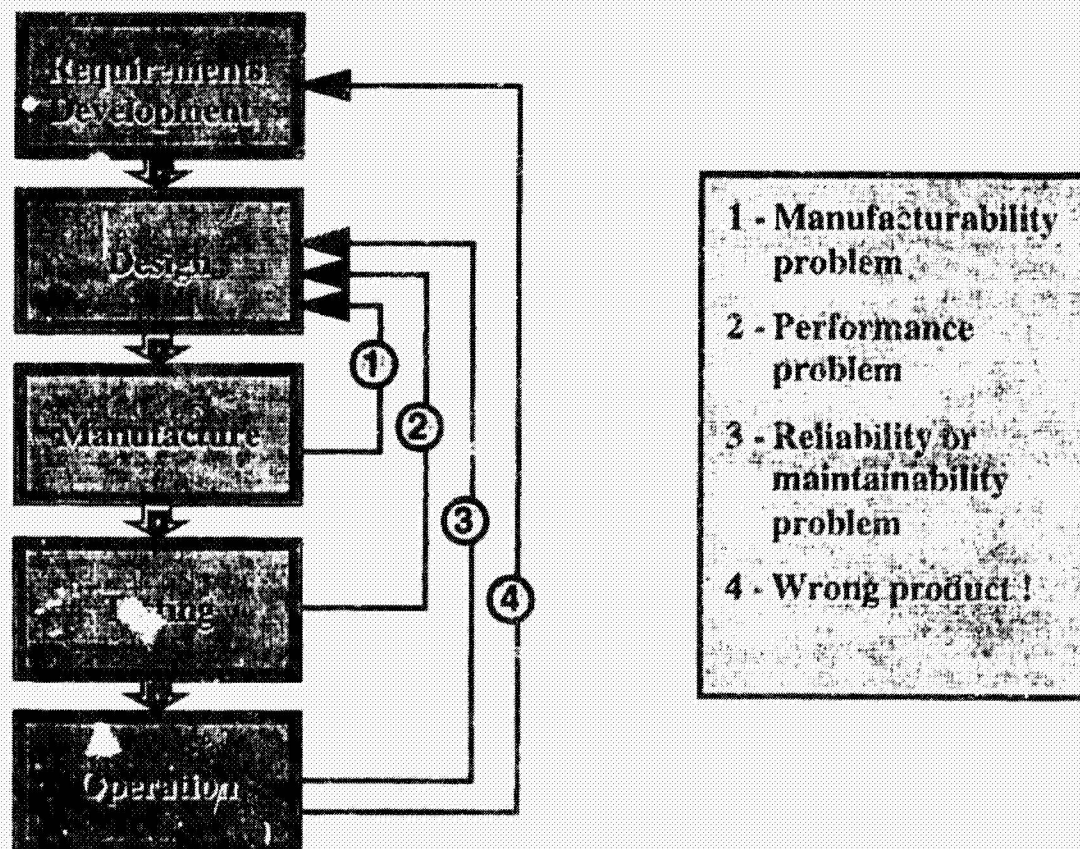


Figure 1 - Linear, sequential product development cycle.

PRODUCT DESIGN CYCLE OF THE FUTURE

As indicated in Fig. 2, the fundamental way of accomplishing the goal of transforming the product development process is to perform product virtual prototyping. Prototyping is not a new concept -- in fact, both NASA and the Department of Defense historically have used this approach for many of their major procurements -- but in the past the prototypes were physical manifestations of the end product. Consequently, such prototypes were themselves costly and time-consuming to create. In synthetic design environments, however, a computer-generated or *Virtual Prototype (VP)* of the product is created. This virtual representation of the design can then be assessed and evaluated, using appropriate simulations of the various phases of the life-cycle. As we shall illustrate in the remainder of the paper, use of virtual prototypes enables true concurrent engineering of a product, which can result in dramatic improvements in quality, innovation, efficiency and customer responsiveness.

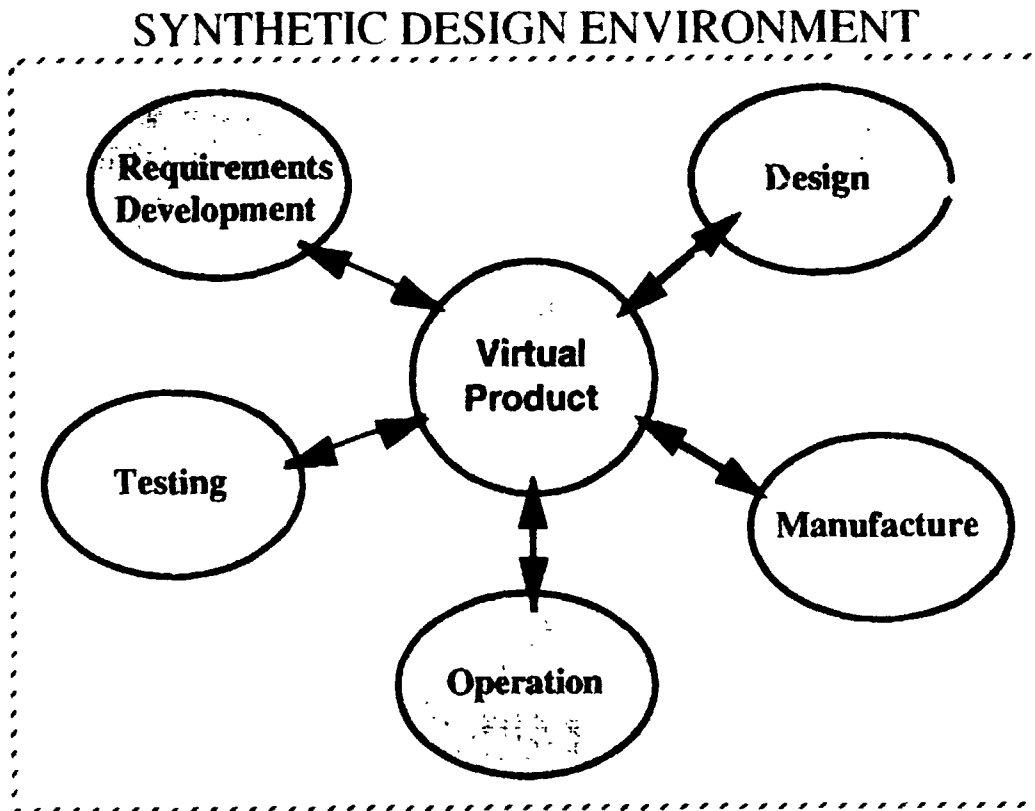


Figure 2 - Future, concurrent product development cycle.

EVOLUTION OF COMPUTER AIDED DESIGN

As outlined in Fig. 3 below, the role of computers in the design process has been evolving over the last thirty years. Initially, they were used by engineers exclusively for the analysis of the design; visualization and design capture were done on drafting tables. These analyses were each individually crafted on the computer codes and were one-of-a-kind products. The first major change was brought about by the introduction of computerized drafting tools, which took the designers away from the drafting tables and put them in front of special-purpose computer display scopes. There the designers captured their designs as electronic drawings on the computer screen by using light pens and special key pads. In parallel, the analysis community was making great strides in creating general purpose analysis codes, e.g. NASTRAN, and computer programs that linked several analyses together were written to perform integrated design and optimization. This was followed by a flurry of development of general purpose codes that became products themselves. These tools harnessed evolving computer technology to perform computer aided engineering, computer integrated manufacturing, and computer aided software (and system) engineering.

In the last decade, computers have become extremely fast at performing calculations and displaying or rendering images. These capabilities have brought about extraordinary advances in the types and levels of analysis that can be undertaken by computers and in the visualization of the results of these analyses. Consequently, computer simulations are now routinely conducted to address extremely complex and detailed physical behavior. Simultaneously, CAD, CAE and CIM tools are being linked together to provide powerful systems for the aiding of the design, engineering and manufacture of products. Added to these advances is the most recent development of synthetic environments, in which designers can immerse to better visualize and understand the product, and in which the products can be "operated" or tested in a simulated real world environment. The confluence of these recent trends is bringing about the next step, the development of synthetic design environments for simulation aided (based) design.

- **Uncoupled Analyses**
- **CAD**
- **Integrated Design (Optimization)**
- **CAE / CIM / CASE**
- **Engineering Simulations**
- **Linked CAD / CAE / CIM**
- **Synthetic Environments**
- **Synthetic Design Environments and Simulation Aided (Based) Design**

Figure 3 - The evolving role of computers in design.

NEXT GENERATION DESIGN SYSTEMS

As noted in Fig. 4 below, we refer to the next generation of computer-based design systems as ***Simulation Based Design (SBD) Systems*** or, alternately, as ***Synthetic Design Environments (SDE)*** .

Simulation based design refers to one central aspect of these new systems -- the emphasis on simulations as a unifying method of representing and analyzing products throughout their life cycle. As used herein, "simulation" has a broad connotation and encompasses modeling/analysis for various purposes and at various levels of fidelity. SBD will utilize and combine virtual, constructive, engineering and physics-based simulations within a single framework to create different views of the product and its behavior throughout its life.

Synthetic design environment highlights another fundamental aspect of the systems -- the creation of artificial or virtual design environments within which diverse representations and models of the product are created, viewed, analyzed and operated. In other words, SDE is the design space for building and conducting product simulations. These simulations may include approximate parametric models for rapid design trades and requirements development; complex, detailed performance prediction models; virtual models for real-time war gaming; manufacturing process models; and virtual reality representations in which customers and designers can examine the product prior to its actual physical creation.

Next generation of computer-based design systems:

- **Simulation Based Design (SBD) Systems**
- **Alternately, Synthetic Design Environments (SDE)**
- **Why Simulation?**
 - **Emphasis on simulations as a unifying method of representing and analyzing products throughout their life cycle.**
- **Why Synthetic?**
 - **Creation of artificial or virtual design environments within which diverse representations and models of the product are created, viewed, analyzed and operated.**

Figure 4 - Next generation design systems will be simulation based, synthetic design environments.

SIMULATION BASED DESIGN

As shown in Fig. 5, the central concept in simulation based design is the creation and use of Virtual Prototypes in synthetic environments to define, develop, produce, test and maintain a complex system. A Virtual Prototype (VP) is a computer generated model of a product capable of functioning properly and appearing realistically in a responsive virtual environment, representative of the physical environment in which the actual product will or may function. Thus, a VP must have the proper geometry and a realistic appearance. It must also respond to the synthetic environment and function in it as it would in the real physical environment.

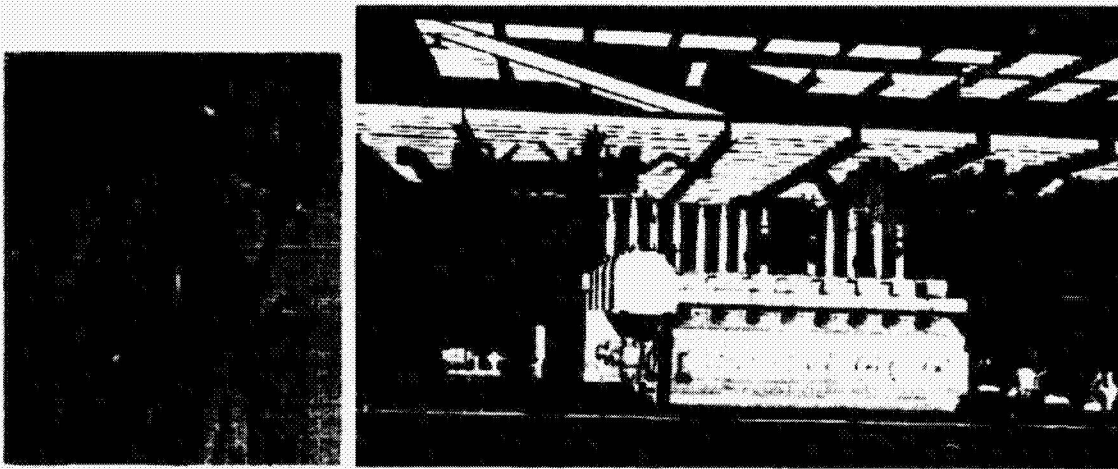


Figure 5 - Virtual prototypes in synthetic environments are central to SBD.

KEY SBD CAPABILITES

Although the constituent capabilities of SDEs are still evolving, several critical features of such systems have emerged. Figure 6 shows five of these key attributes: (1) an immersive design environment to facilitate collaborative product viewing and interrogation; (2) high fidelity computerized, physics-based simulations to properly model the performance of the product; (3) distributed interactive simulations to enable the product to be tested or operated in a synthetic environment; (4) a visual programming environment in which engineering and manufacturing analyses can be rapidly linked together; and (5) an integrated, "smart" product model that captures the data, models and analyses needed to represent the design and the processes by which it is produced.

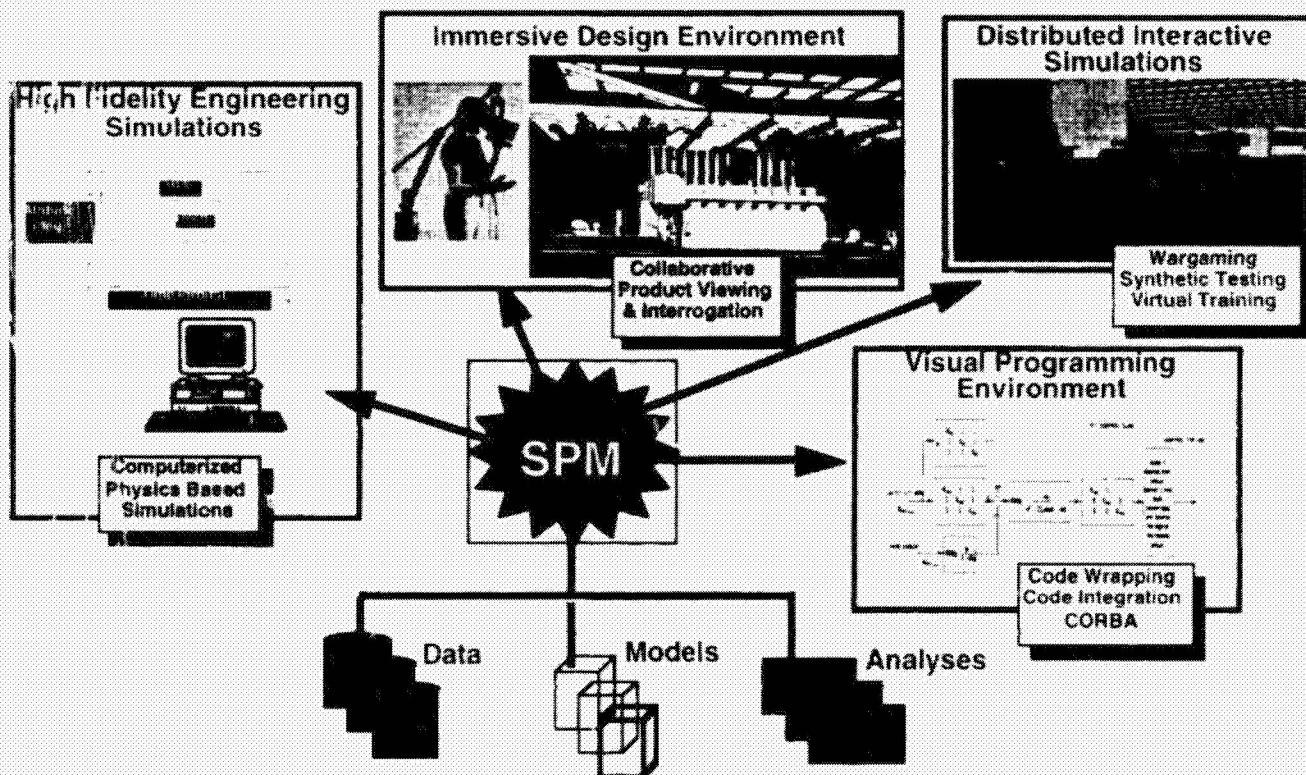


Figure 6 - Integration of key SBD attributes.

NETWORKING EXPERTS

Implicit in the design environments of the future is the concept illustrated in Fig. 7; i.e. that the environment will support design teams that are geographically distributed. Modern computer network technology enables us to connect design team members who are not physically collocated. Similarly, the technology enables the design itself to be physically distributed among data sets residing on computer hardware that is not collocated. Typically, even the analyses codes can and will be located on different machines in different physical locations.

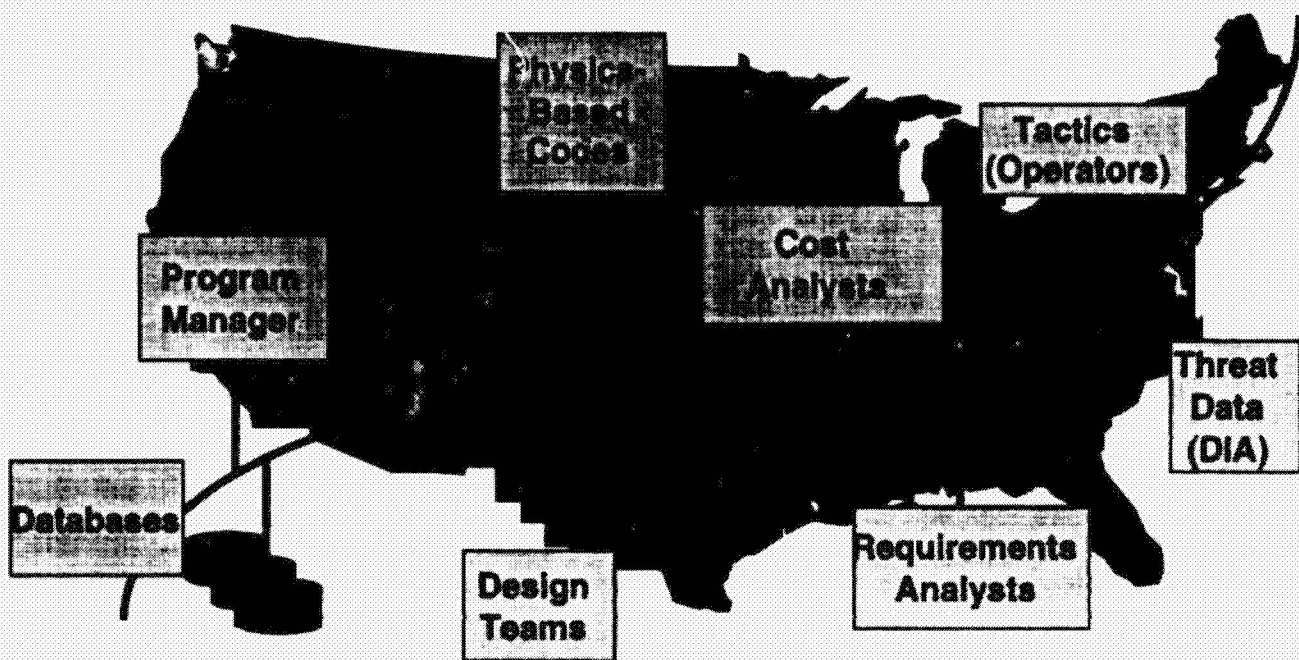


Figure 7 - SBD will tie together geographically distributed design teams.

DEMONSTRATION OF SIMULATION BASED DESIGN

An intensive demonstration of how an advanced synthetic design environment might work in the shipbuilding industry was conducted in the summer of 1995 by two teams funded by the Advanced Research Projects Agency. The remainder of this paper addresses the demonstration conducted by a team from Lockheed Corporation, Science Applications International Corporation and Newport News Shipbuilding. The design environment was given the name, Simulation Based Design (SBD). As shown in Fig. 8 below, SBD was used to conduct a design exercise for a notional military roll-on/roll-off ship (denoted the Notional Baseline Ship, NBS), having amphibious lift capability.

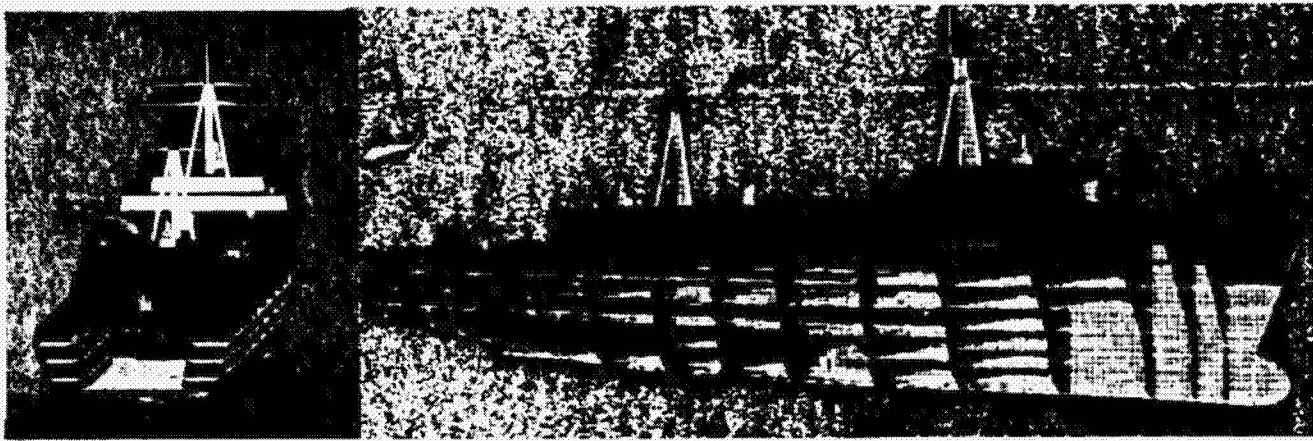


Figure 8 - The Notional Baseline Ship, NBS.

SHIP DESIGN DEMONSTRATION SCENARIO

As illustrated in Fig. 9 below, the ship design demonstration scenario consisted of seven elements or stages: (1) mission analysis, (2) propulsion system selection (featuring use of the smart product model), (3) collaborative design, (4) distributed interactive simulation, (5) multi-disciplinary analysis, (6) manufacturing analysis, and (7) cost and risk analysis.

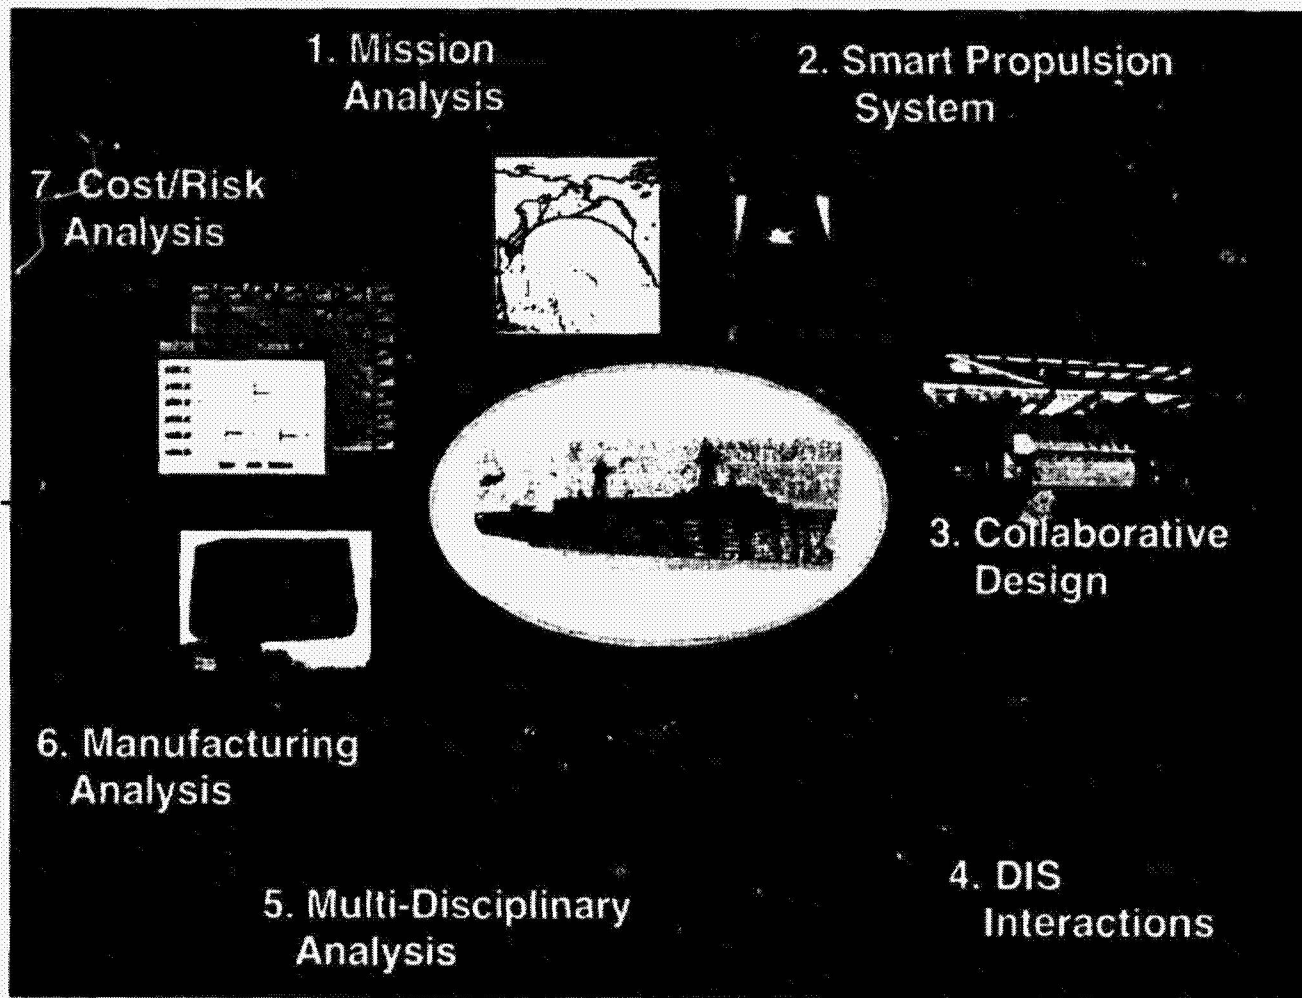


Figure 9 - Segments of the SBD ship design demonstration.

MISSION ANALYSIS SEGMENT

In the first segment, data on the existing NBS design are used in an operational analysis to determine the required speed change and the associated engine horsepower change. Concurrently, a first estimate is made of the associated cost. As seen in Fig. 10, this segment begins with a hypothetical mission requirement to deliver men and material from ports in the continental United States to the Korean peninsula within a prescribed time period. A presumed fleet of existing NBS must be evaluated to determine the ability to perform this mission.

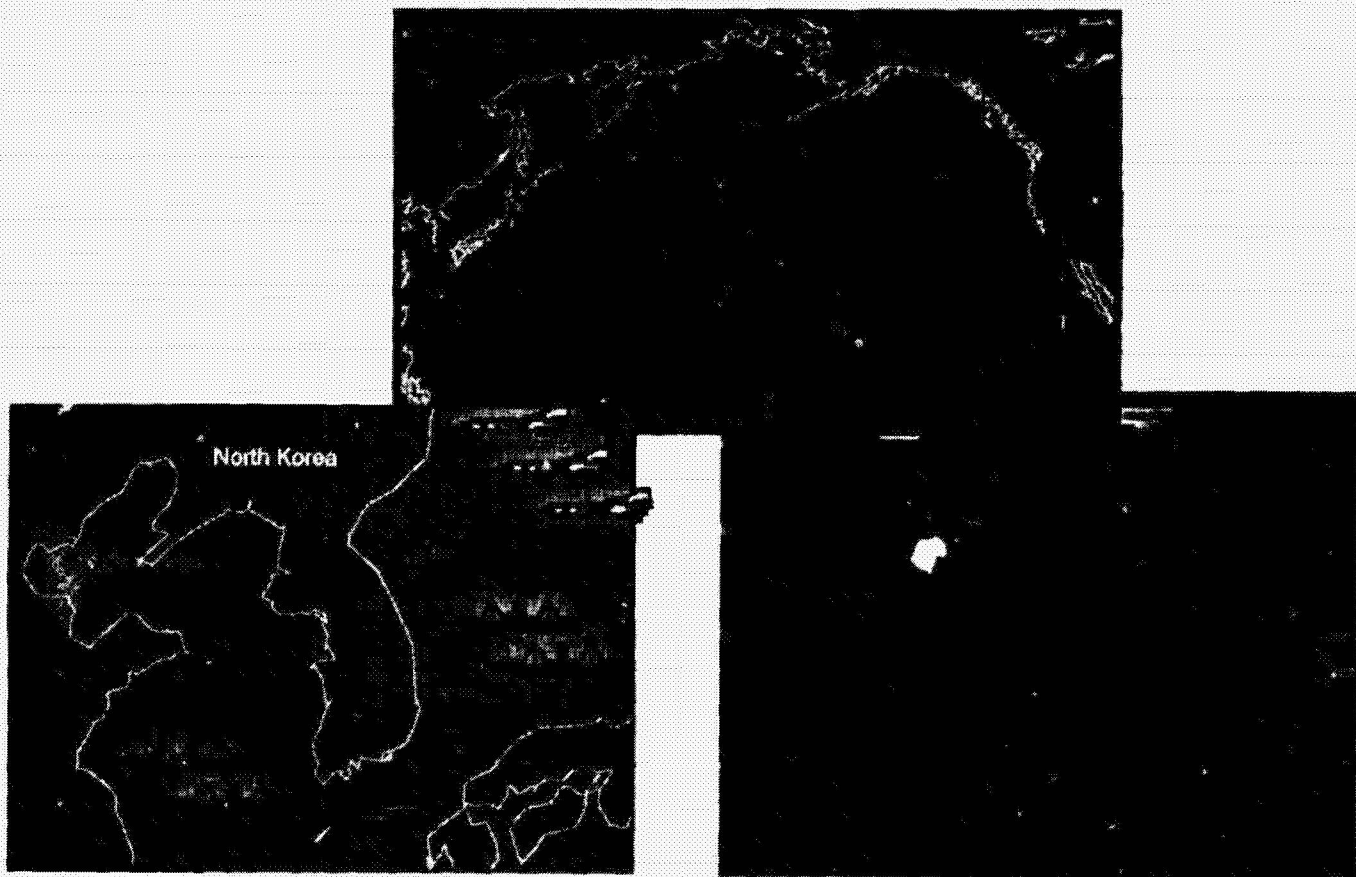
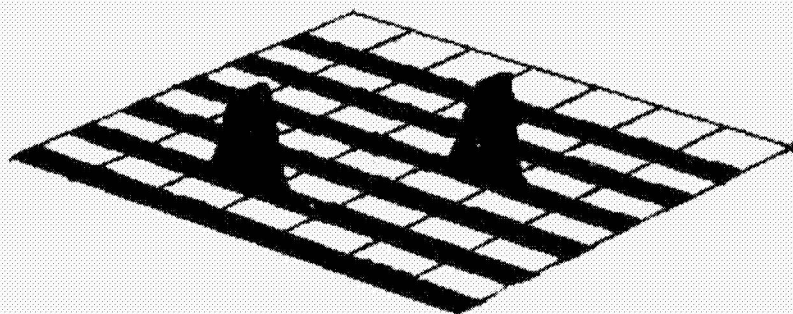
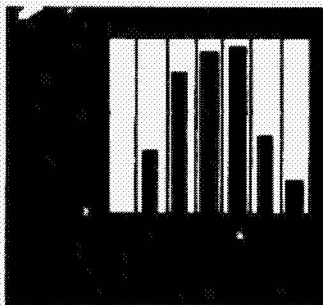
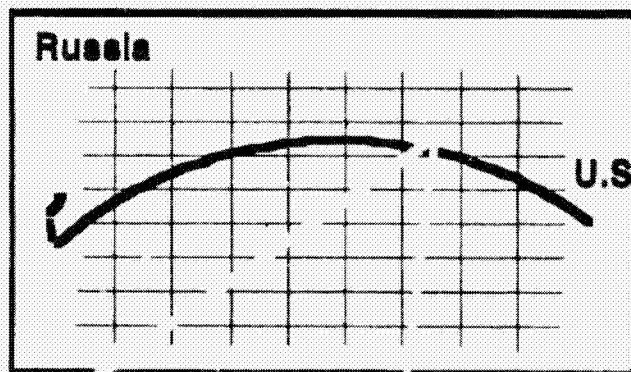


Figure 10 - Hypothetical mission for the NBS.

MISSION ANALYSIS USING SPEED-OVER-GROUND MODEL

Using existing data on the NBS, SBD performs a rapid mission analysis. As shown in Fig. 11, Monte Carlo simulations of the ships' transit are conducted, using statistical data on sea state along the specified route and NBS resistance curves, derived from hydrodynamic analyses. The data and models for this speed-over-ground (water) analysis are "stored" in (i.e., associated with) the product model for the NBS. This product model is referred to as a "Smart" product model, in part, because such associations can be made to assure that appropriate data and analyses are automatically accessed whenever the design is interrogated.



- ❑ **Integrates SPM with Operations Model**
- ❑ **Exemplifies Use of an Operational Model in the Design Process**

Figure 11 - Monte Carlo analyses of the mission are performed using data and models stored in the Smart Product Model for the NBS.

MISSION ANALYSIS - ALTERNATIVES FOR INVESTIGATION

After the initial analysis shows that the existing fleet cannot meet the required material build-up, several alternatives are postulated for further consideration. Two of these alternatives -- increase the number of NBS in the fleet (i.e., build more ships), and refit the NBS to increase its speed -- are carried forward for more detailed analysis. Figure 12 shows the results of those analyses: three additional ships would be needed to satisfy the requirement; or the existing ten ships would have to increase their top speed from nineteen to twenty-two knots. The mission analysis also determines that an engine shaft horsepower of 45,000 is required to propel the NBS at the top speed of twenty-two knots.

- **Number of Ships**
- **Speed**
- **Capacity**
- **Decommission Fewer Ships**

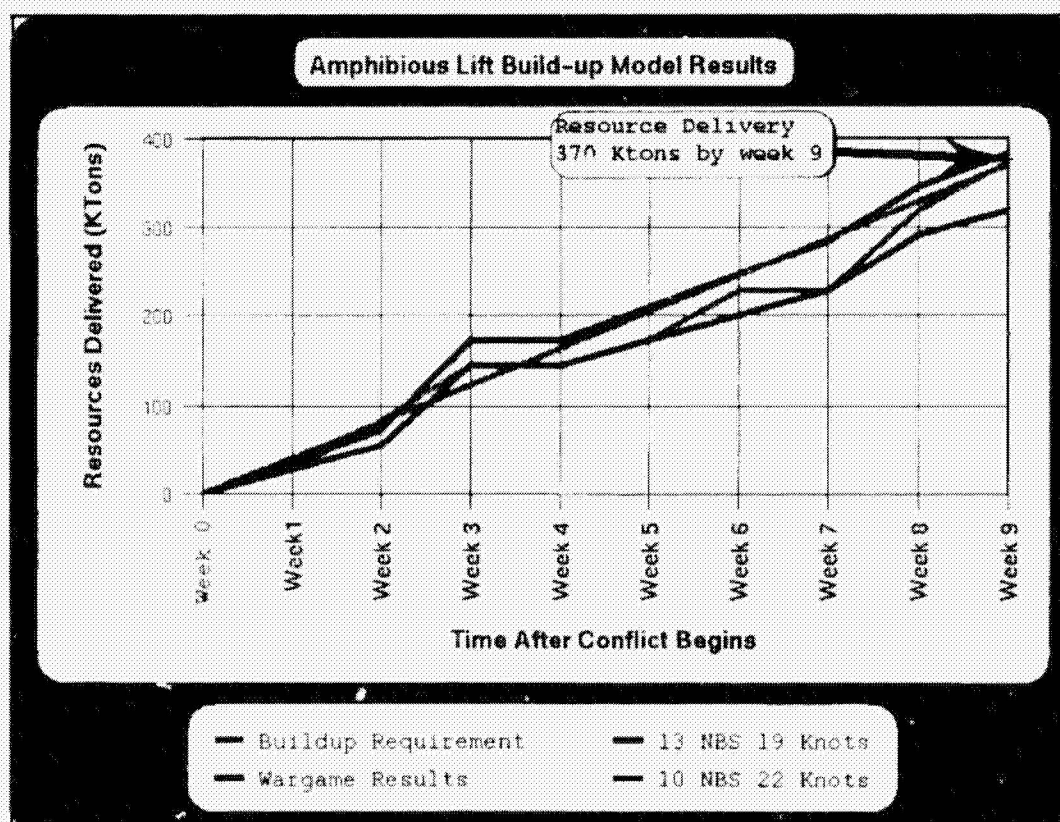


Figure 12 - Results of SBD Mission Analysis.

MISSION ANALYSIS - COST AND RISK ANALYSIS

For the two chosen "solutions," SBD is used to generate a rapid cost and risk analysis. The necessary data are automatically gathered from the Smart Product Model. The associated risk analysis code determines the expected cost of each alternative and the distribution about that expected value. As shown in Fig. 13, the alternative of increasing the speed of the NBS appears to be better than the building of more ships, and appears to be only slightly more costly than the baseline. The remainder of the demonstration explores the faster-ship alternative in more detail. In effect, an actual redesign is performed to confirm the viability of increasing ship speed, explore the impact of the design change on other ship performance attributes, determine the impact on manufacturing and develop a detailed cost estimate to support or refute the conclusion of the preliminary cost/risk analysis.

- **Uncertainty/Risk Management**
- **Recognizes SPM Dependencies**
- **Actuates Engineering Models**
- **Tracks Design Change Impacts**

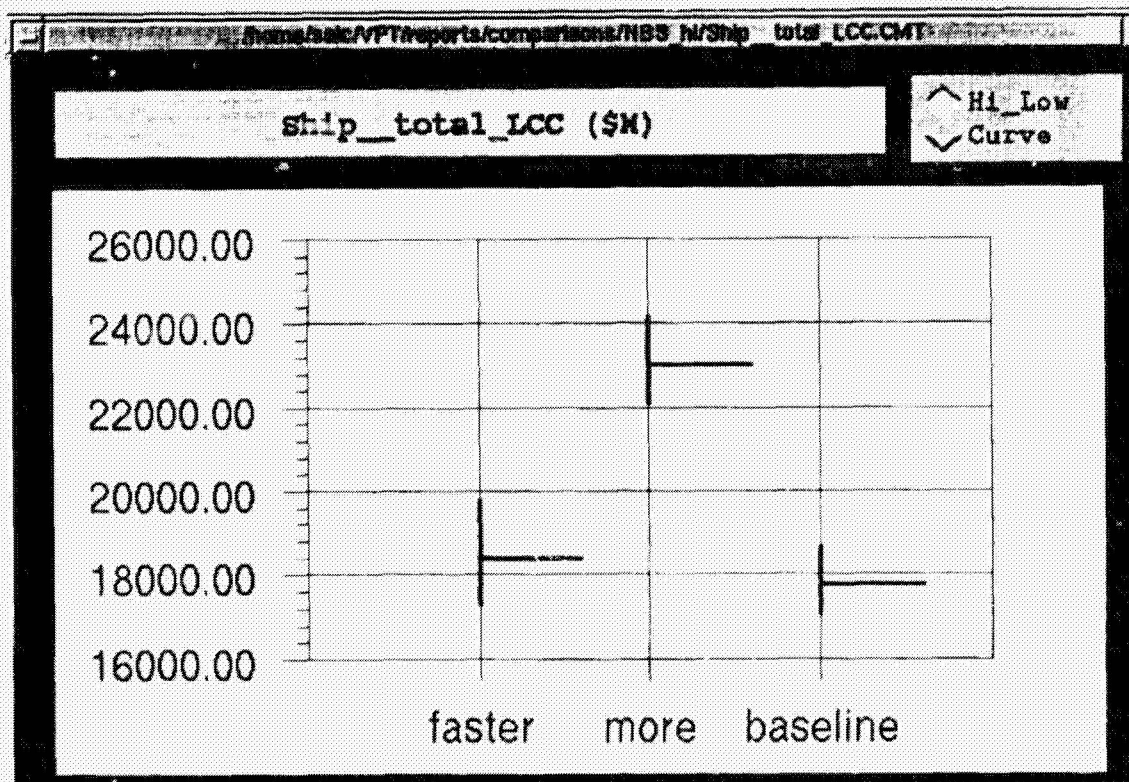


Figure 13 - Cost and risk analysis within SBD is the key to making informed design decisions.

SMART PRODUCT MODEL

In the industry of the future, a model of a product will be constructed at the instance of its conception. This model will be expanded, refined and invoked throughout the product's life to capture and evaluate the evolving design. This model is sometimes called an Integrated Product and Process Development Model (IPPD) because it contains descriptions of both the product and the processes required to develop and produce it. We call the model a "Smart Product Model" (SPM) to denote that higher level product information is captured; i.e., it includes not only a physical description of the product but also the data, models and analyses necessary to fully characterize and evaluate the product. The model is also "smart" in that it knows where, among a myriad of voluminous and geographically distributed electronic databases, to find the data it needs to perform and display the results of any analysis or simulation called for at any given time. The SPM is "smart" in a third sense; it is structured as a collection of objects, each of which has its own set of associated attributes, behaviors and characteristics. For example, if a sub-component of the system/product is replaced, performance and cost models associated with the new and old components are automatically replaced as well. As shown in Fig. 14, the SPM unifies product and process information shared within the enterprise, and acts as a real-time, distributed spreadsheet providing and controlling product and process information across the enterprise throughout the product life cycle.

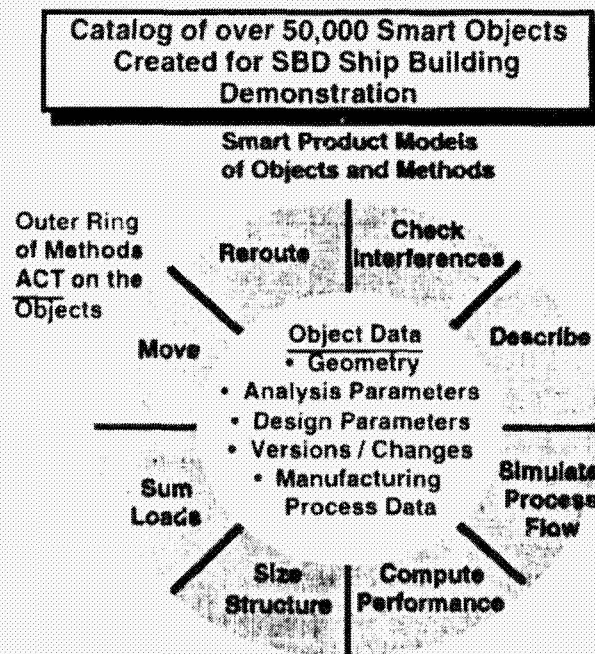


Figure 14 - The Smart Product Model links data and processes in an object-oriented structure.

SMART PRODUCT MODEL - CONTINUED

As shown in Fig. 15, the product is represented in a logical hierarchy in the Smart Product Model. Thus, in the demonstration the ship is an object consisting of other objects, such as the bridge and the engine room. In turn, each of these objects consists of other objects, e.g., the engine room contains an engine, a heat exchanger, etc. In the SBD, one can change the design by altering aspects or features of the constituent objects or by completely replacing old objects with new. The user interface makes the swapping of objects as easy as a "drag and drop" operation on a desk-top personal computer.

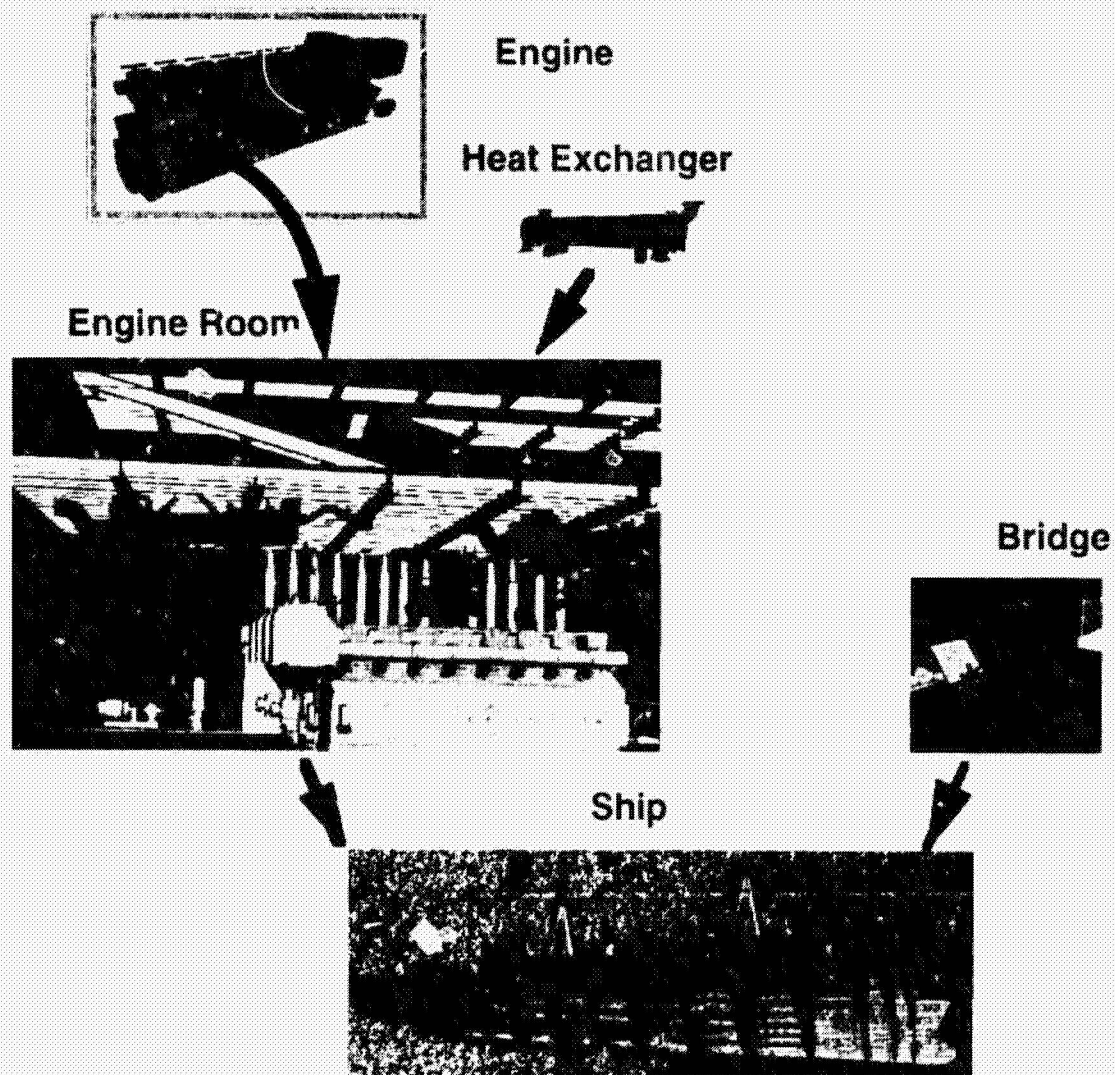


Figure 15 - Drag and drop hierarchy of smart objects populate the Smart Product Model.

PROPULSION SYSTEM SELECTION

From this point forward in the demonstration, the synthetic design environment is used to visually design changes and to examine their impact. This segment employs electronic commerce via the Internet to select an engine replacement capable of providing the necessary horsepower. The product model data associated with the engine are brought into the Smart Product Model (SPM) on the complete NBS with a simple "click and drag" operation, and the new engine automatically appears in the SDE with associated piping automatically re-routed. The SPM automatically invokes and executes associated analyses that identify and flag in the SDE any remaining equipment fouls. Other analyses are automatically executed that determine if the capacity of an existing heat exchanger is inadequate. This leads, in turn, to selection of a replacement for the heat exchanger using electronic commerce, the SPM and the SDE.

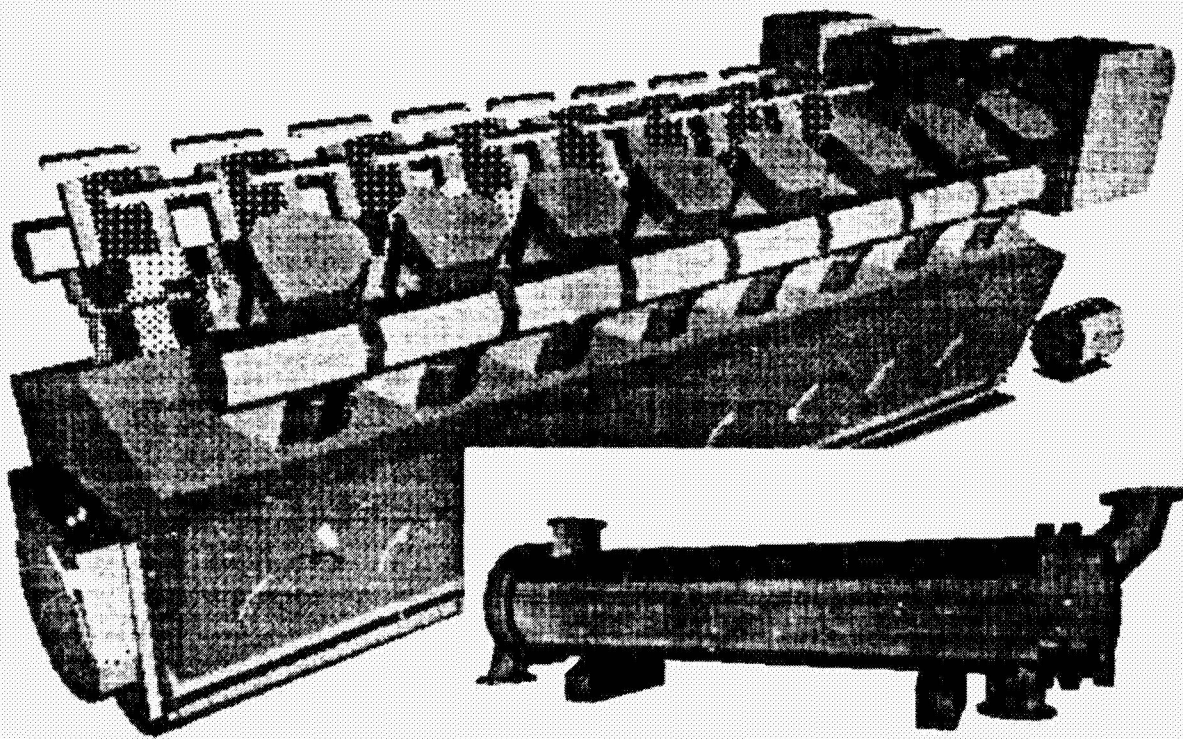


Figure 16 - Engine and heat exchanger models are selected via Commerce Net and swapped into SPM.

COLLABORATIVE DESIGN

In this segment of the demonstration scenario, a Propulsion Engineer and an Arrangement Specialist are simultaneously immersed in the Virtual Design Environment (VDE) to resolve the fouls caused by placement of the new engine and heat exchanger. Both members of the design team "enter" the engine room shown in Fig. 17. One participant uses a FakeSpace BOO3C, stereo, high-resolution color visualization system while the other uses a Virtual Research Eyegen3 Head Mounted Device. Using a three-dimensional pointer, the Arrangement Specialist selects and moves a bulkhead, associated equipment and deck grating to accommodate the larger engine. Piping re-routing is performed automatically while other arrangement changes are commanded directly by the immersed design team. By obtaining virtually instantaneous three-dimensional design updates, the team is able to quickly make arrangement changes, and investigate and verify impacts. One additional feature illustrated in this segment is automatic notification of other, non-immersed members of the design team that a design change was made in an area that might impact them.

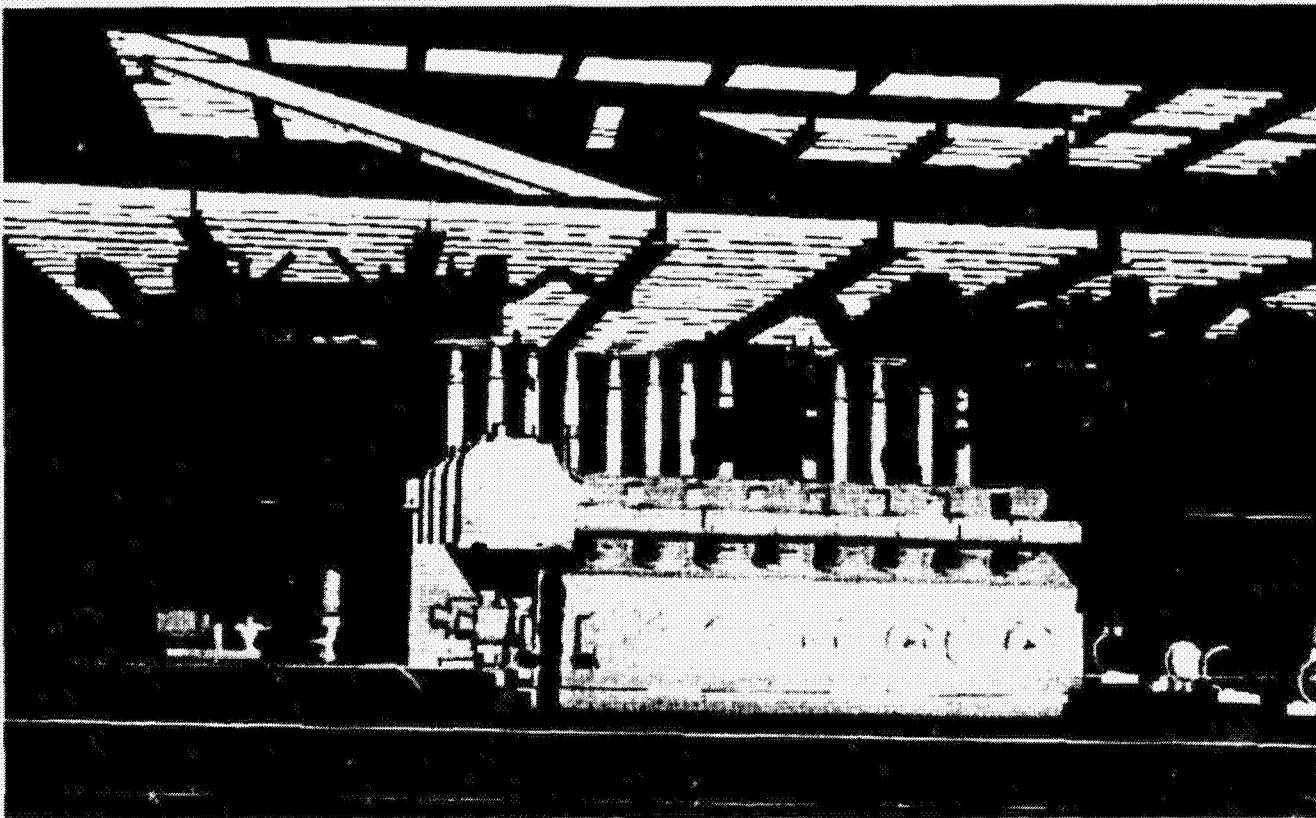


Figure 17 - Engine room as displayed in the Virtual Design Environment.

DIS INTERACTIONS

In response to the design change notification issued in segment 3, the designer responsible for the stowage compartment above the engine room is alerted to a structural change that reduces the volume of the stowage compartment in certain areas. To determine whether this change compromises the ability to load armored tanks, a Distributed Interactive Simulation is invoked wherein an M1 tank driver attempts to drive his vehicle aboard the NBS and into the stowage compartment. In Fig. 18, the tank is shown driving up a ramp in the interior of the NBS. The unique features of this demonstration are: (1) The driver and his simulator are located in a physically different place from the rest of the participants so that his entire view of the ship is created by PDUs sent from the ship VDE. He is given only an out-of-the-window view from his tank. (2) The tank's movements are transmitted back as DIS Entity-State data via PDUs and are displayed within the ship VDE. The rest of the design team has freedom to view the scenario from any and multiple design perspectives. (3) Both the tank and the NBS are virtual prototypes created in their own separate SDE but joined together in this simulation. As the tank driver maneuvers his vehicle, any fouls between the tank and the ship are automatically flagged and recorded for the entire design team to evaluate.

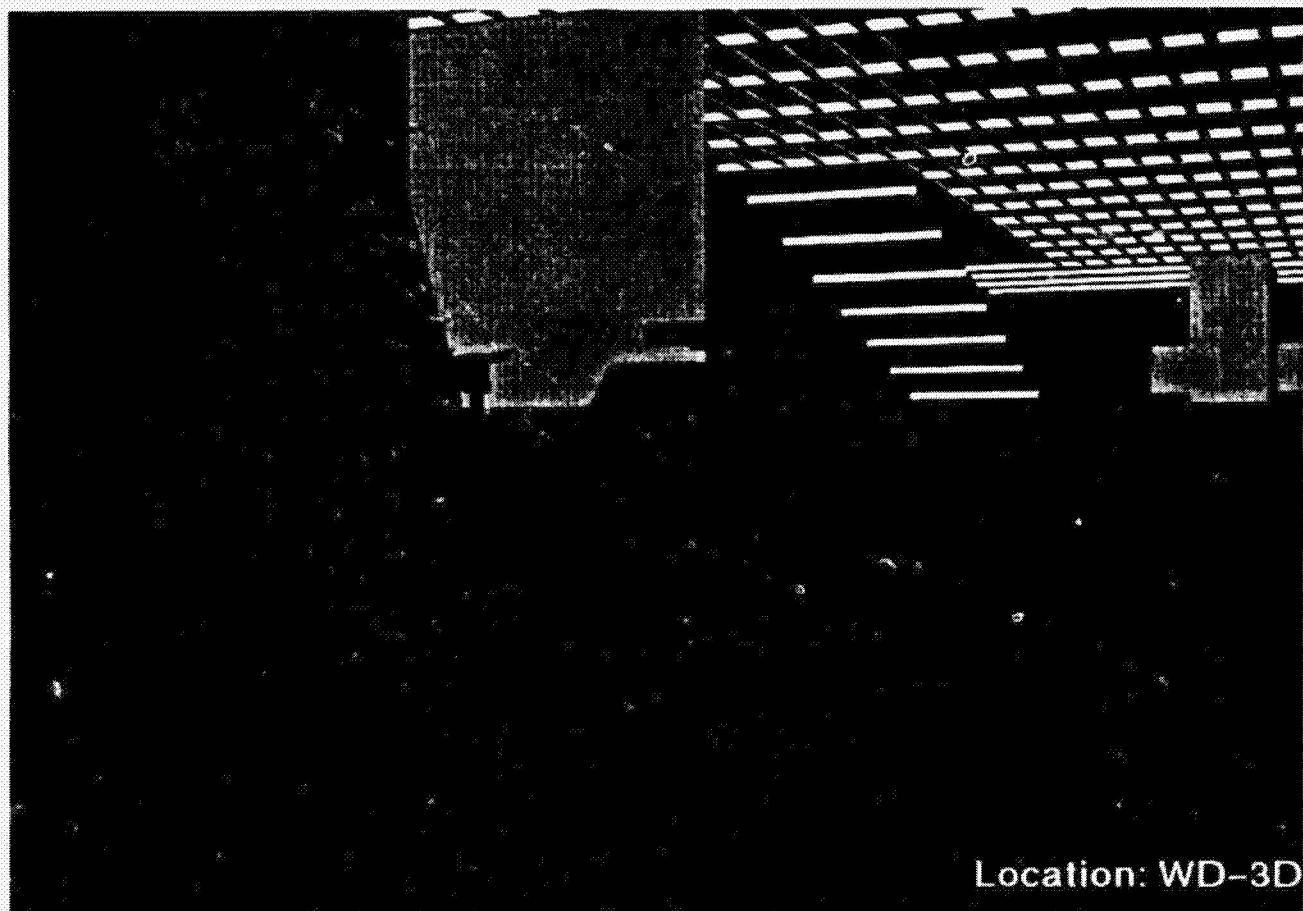


Figure 18 - Tank simulation is linked to the NBS in SBD's Virtual Design Environment.

MULTI-DISCIPLINARY ANALYSES

In this segment, complex engineering analyses necessitated by the design changes are planned and executed. Due to the increase in ship top speed (new engine), together with the weight and structural arrangement changes, a new computation of dynamic loads resulting from ship response in the seaway is required. Using visual programming tools, an analysis procedure is set up that uses hydrodynamic codes, coupled to linear ship response codes, to compute ship motion for a range of ship speed, heading and sea states. As shown in Fig. 19, the resulting ship motion is displayed, using VDE and the viewer sees the NBS traversing a stressing seaway condition. From examination of these results, the ship responses are judged to be too large to base reliable load predictions on linear theories. A separate nonlinear analysis of critical conditions is then set up and routed to a remote supercomputing facility to be executed. External loads from this analysis are transferred to a FEM of the NBS and internal loads are computed. Again using VDE, members of the design team immerse themselves within the ship to examine stress concentrations in areas of the design changes. This examination reveals excessive stress in the deck, leading the structural analyst to decrease stiffener spacing via the VDE to solve the problem.

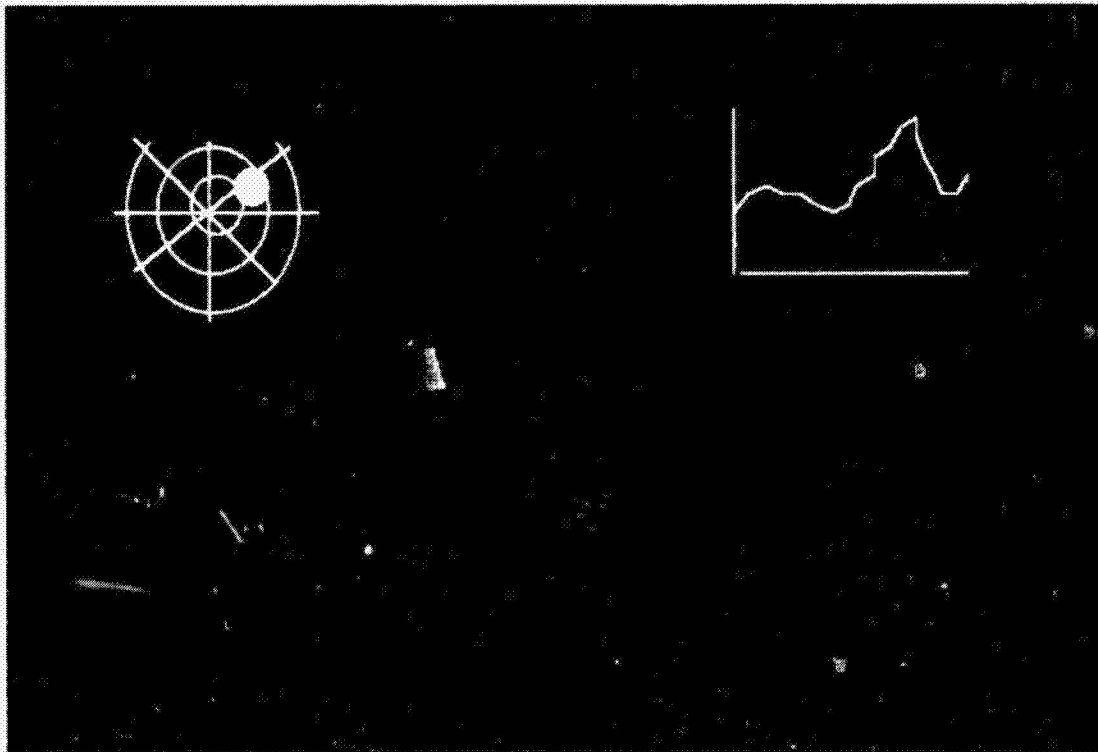


Figure 19 - Ship's motion and dynamic loads are viewed in SBD's Virtual Design Environment.

MANUFACTURING ANALYSIS

Notified of all the design changes made in the previous segments, a Manufacturing Specialist invokes his model of the manufacturing processes required to construct the NBS. As indicated in Fig. 20, all design change data are automatically incorporated in this analysis through its linkage back to the Smart Product Model. For this exercise, the G2 discrete event simulation package from Gensym is used to model the work flow processes. The segment focuses on simulating activity within the shipyard's steel fabrication facility. Cost estimates of each process are generated and accumulated as the simulation executes. To lower total cost, a process modification (addition of a second blast and coat operation) is made and is found to reduce schedule, inventory costs and, hence, total cost.

- Gensym's G2 discrete event simulation
- Smart Product Model automatically linked updated design, process and cost data / models to simulation

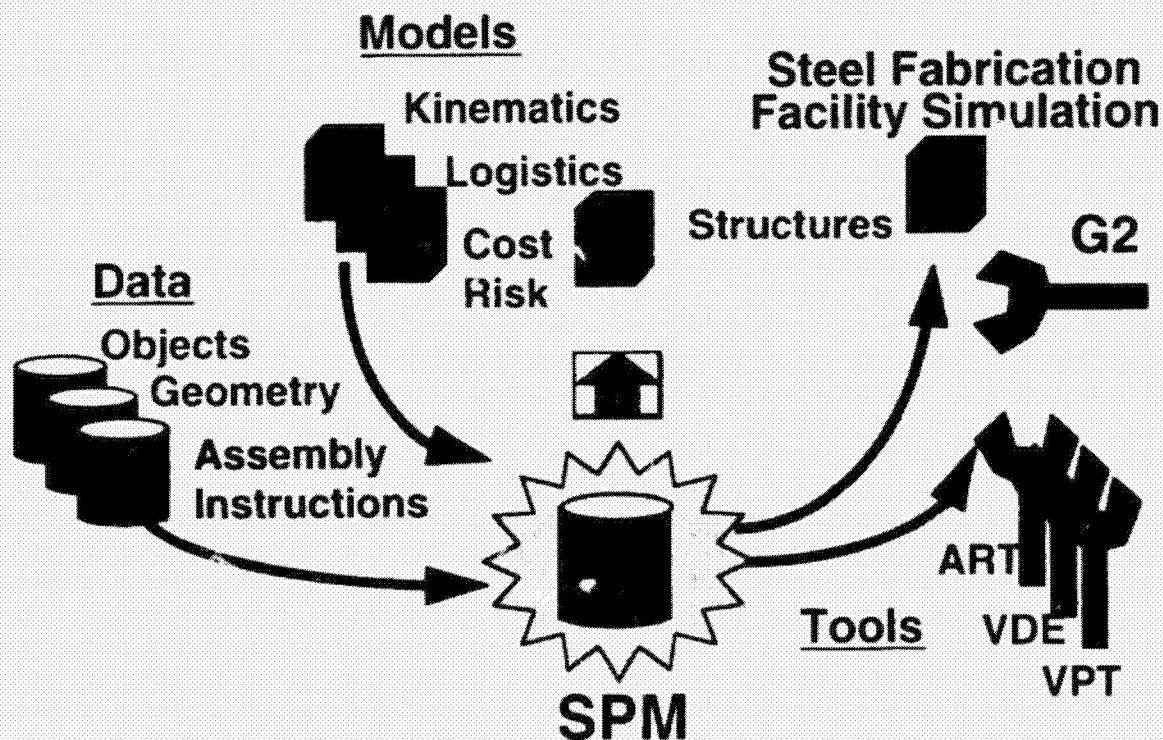


Figure 20 - SBD's Smart Product Model links design changes to manufacturing process simulation.

COST/RISK ANALYSIS

When the scenario originally began, a cost estimate was made using historical data available from the baseline NBS design. During the various segments, as design and process changes were made, data that have a direct bearing on the cost were changed. SBD includes a tool (VPT) that keeps track of the cost impacts of the changes as the design modification progresses. As shown in Fig. 21, VPT provides dynamic costing information, in effect creating a "cost meter" to help identify cost problems as they arise and evaluate solutions as they are proposed.

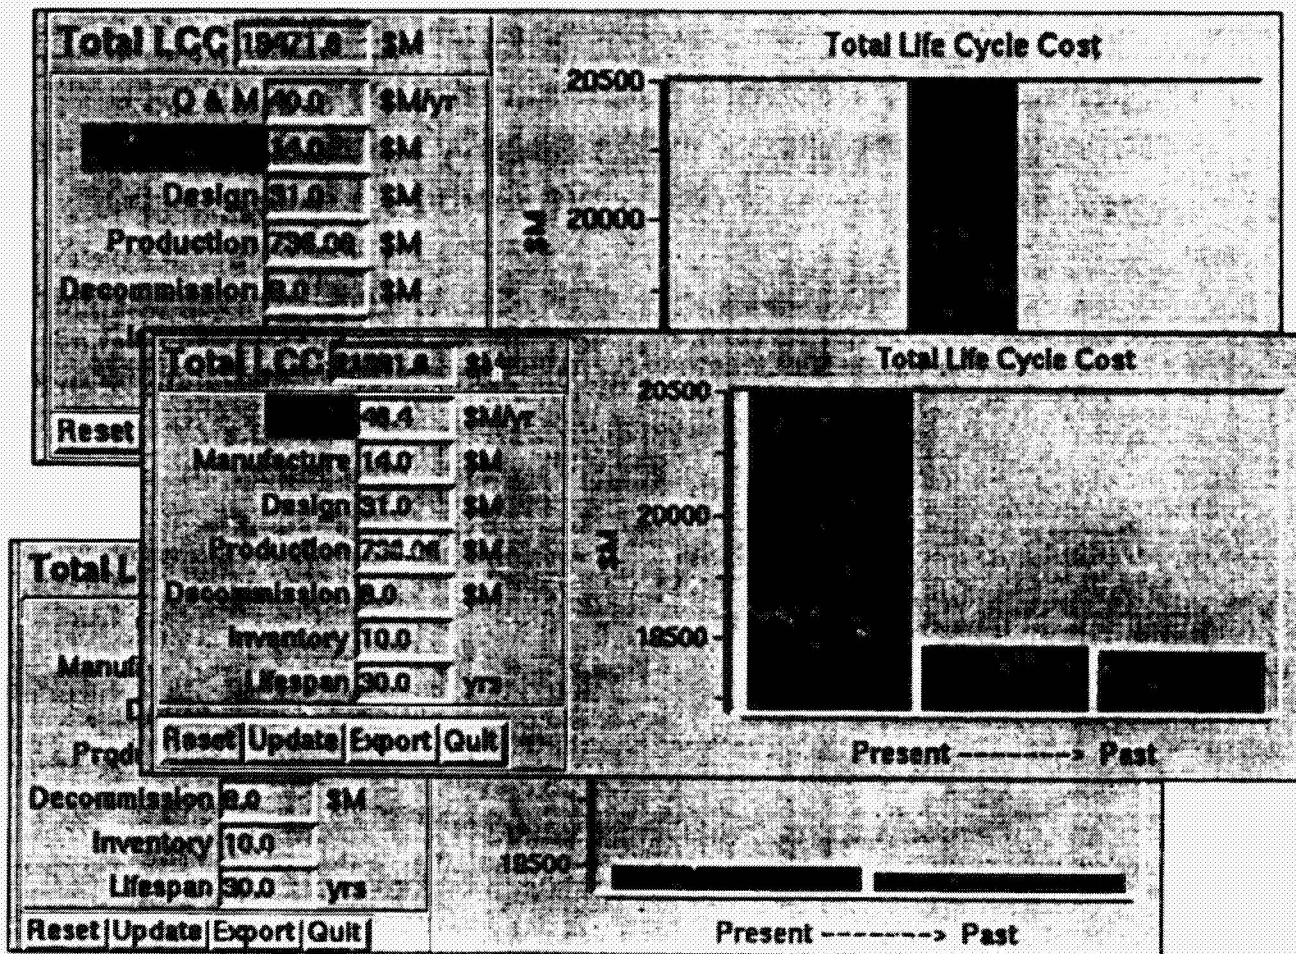


Figure 21 - SBD's Virtual Prototyping Tool (VPT) provides dynamic costing.

COST/RISK ANALYSIS (CONTINUED)

At the end of the design exercise, final cost and risk analyses are performed to definitize the likely impact of the changes, including the ever-present uncertainty in cost estimates. As shown in Fig. 22, these results show that the final refined analysis has less uncertainty in its predicted cost and that the predicted (most likely) cost is similar to the original estimate. Thus, these analyses produce results that confirm the correctness of the original decision to modify the design to increase ship top speed, rather than simply increasing the production run.

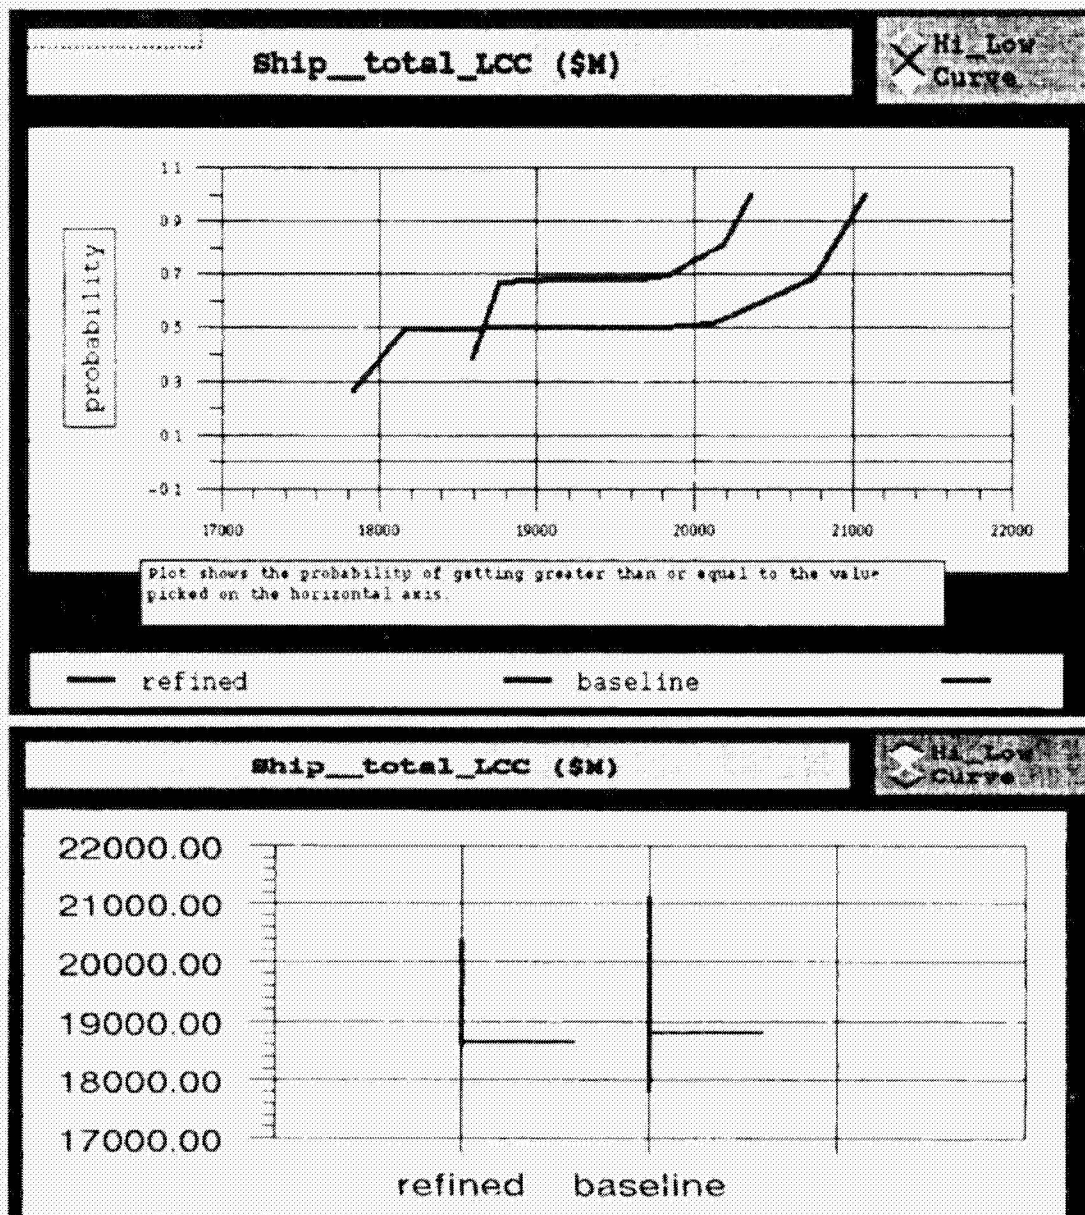


Figure 22 - SBD's VPT computes uncertainty in basic design metrics such as cost.

CONCLUSION

The series of analyses and design activity performed in the demonstration scenario were representative of a set of activities necessary to perform the required design study. While the demonstration took approximately three hours to complete, performance of a similar scope activity without SBD technology was estimated to require two months -- a two-order-of-magnitude difference. As indicated in Fig. 23 below, the SBD demonstration convinces us that practical, powerful virtual design environments are feasible today.

- **Shipbuilding demonstration showed a two-order of magnitude improvement in design cycle response time**
- **SBD demonstration showed that a powerful VDE is feasible today**
- **Necessary technologies are available and improving rapidly**
- **Future advanced immersive environments will lower cost and improve realism**
- **Development of a prototype, general-purpose virtual design environment (i.e., SBD) is timely**

Figure 23 - Conclusions of SBD demonstration.

SUMMARY

Figure 24 summarizes this paper. Advances in computer hardware, computer software and computer science have provided the tools needed to revolutionize the traditional product development cycle. The evolution of computer aided design has reached the age of synthetic design environments, in which computer-generated virtual prototypes of complex systems will be the central unifying basis for all product development activities. Design, engineering, manufacture planning, product test, and operator training will all be performed with virtual prototypes in synthetic design environments (SDE).

To demonstrate the feasibility of this bold concept, a notional SDE was constructed for the ship building industry. Denoted as the Simulation Based Design (SBD) environment, this SDE was used to conduct a demonstration ship design scenario. Starting from the definition of mission requirements and following into design, engineering and manufacture planning, this demonstration used many advanced technologies to show how an SDE might operate and to quantify the improvement that it could produce in the design cycle of a product as complex as a modern ship.

The integration of the technologies embodied in SBD was shown to lead to a large (two orders of magnitude) productivity increase in the product development process. Central ingredients of SBD were the immersive design visualization environment, the "smart" product model, the integrated cost and risk tool, and the visual programming environment. SBD offered a rich human-computer interface to create a design space in which a team of engineers and designers viewed and interrogated the product as they designed and analyzed it. The precise nature and quality of the human-computer interface provided by future synthetic design environments will determine their acceptance, utility and, ultimately, the extent of true productivity gains.

- **Advances in computer technology and computer science are revolutionizing the design process**
- **Next step in evolution of computer aided design is the Synthetic Design Environment (SDE)**
- **An SDE, known as Simulation Based Design (SBD), was demonstrated for shipbuilding**
 - **Immersive visualization**
 - **Smart product model**
 - **Integrated cost and risk assessment**
 - **Visual programming environment**
- **Acceleration of design process by two orders of magnitude is possible**

Figure 24 - Summary of the demonstration of a synthetic design environment for shipbuilding.

**NEXT
DOCUMENT**

S11-62
0250

DISTRIBUTED AND COLLABORATIVE SYNTHETIC ENVIRONMENTS

Chandrajit L. Bajaj and Fausto Bernardini
Department of Computer Science
Purdue University
West Lafayette, IN 47907
{bajaj, fxb}@cs.purdue.edu

12/15/11

Introduction and Overview

Fast graphics workstations and increased computing power, together with improved interface technologies, have created new and diverse possibilities for developing and interacting with synthetic environments (Refs. 1-6). A synthetic environment system is generally characterized by the following components:

- Input/output devices that constitute the interface between the human senses and the synthetic environment generated by the computer. Several degrees of immersion are possible, ranging from simple stereoscopic view of an image on a CRT display to a total immersion in which a head-mounted display, sound and haptic devices (force and torque feedback, tactile stimuli) are used.
- A computation system running a real-time simulation of the environment. This can sometimes be subdivided in several subsystems; for example, a simulation module running on a supercomputer coupled with a scene creation and a rendering module running on a graphics workstation.

To achieve an acceptable level of realism, the display subsystem must generate at least ten frames per second of high quality graphics. For complex environments, this goal is still far from being reached. Nonetheless, synthetic environments have already been applied successfully in such diverse fields as: operations in hazardous or remote environments, through telepresence; molecular modeling; flight simulations; battlefield simulation; architectural walk-throughs; surgical planning; collaborative design; education and training; entertainment; and many others.

A basic need of a synthetic environment system is that of giving to the user a plausible reproduction of the visual aspect of the objects with which he is interacting. It is well known that populating a synthetic world with objects, whether they be spaceships in a science fiction movie, prototypes in a manufacturing design, or replicas of existing real-world objects in an architectural walk-through, can be extremely time consuming. Not only must it be possible for the user to move in the environment, but updating in real time his or her view of it. To convey the impression of an immersive, active presence in an environment, a real-time simulation of the physical behavior of the various components of the environment is necessary. Moreover, in many applications the user should be allowed to interrogate objects about their associated properties, and to interact with them; for example, modifying some of their properties to support *what-if* scenarios.

The goal of our Shastra research project is to provide a substrate of geometric data structures and algorithms which allow the distributed construction and modification of the environment, efficient querying of objects attributes, collaborative interaction with the environment, fast computation of collision detection and visibility information for efficient dynamic simulation and real-time scene display. In particular, we address the following issues:

- A geometric framework for modeling and visualizing synthetic environments and interacting with them. We highlight the functions required for the *geometric engine* of a synthetic environment system.

- A *distribution and collaboration* substrate that supports construction, modification and interaction with synthetic environments on networked desktop machines.

Geometric Engine

A critical subsystem in all synthetic environment systems is the *geometric engine*, or the software module responsible for creating a realistic view of the simulated world, and for allowing user interaction. In a typical scenario, a system requires the display of several objects that move and must behave realistically. A user wanders in the environment, constantly changing his or her point of view. Other users (or actors) may be sharing the same environment, and a suitable representation of them could be required. The users interact with objects in the environment; for example, touching their surface to pick and query an object, grabbing and moving them, or simply colliding with a wall of a room. Simulation of even the simplest form of dynamic and interaction requires collision detection and contact analysis. Fast display of complex environments requires efficient visibility computation. Querying and interacting with objects requires rapid point location and local shape control. In short, the geometric engine of a general-purpose synthetic environment system should provide efficient data structures and algorithms for:

- shape representation
- dynamic object insertion and deletion
- object animation (motion, non-rigid transformation)
- object location and closest point queries
- collision detection and contact analysis
- visibility ordering and culling
- view-volume clipping
- multi-resolution representation
- real-time high-quality rendering
- other operations required by specific applications (e.g., set operation and interactive shape control for collaborative design).

We want to explore the use of data structures suitable to support the operations listed above in a complex, constantly changing environment. Several types of data structures that partially satisfy these needs have been proposed in the past. However, they have mainly been used in Computer-Aided Design (CAD) systems or other special-purpose applications, and not as an infrastructure for a complex, dynamic and general-purpose environment. Often, CAD systems use boundary representations (Breps) to describe the geometry of the object being modeled. Breps are well suited to the implementation of the most common modeling operations. However, when this representation is to be used in a general-purpose synthetic environment, it has to be supplemented with additional structures such as octrees or bounding volume hierarchies to achieve the required efficiency.

Several variants of the *octree* data structure have been proposed, and be used either as a superimposed search index on existing representations of geometry, or as the main representation scheme in the system (Ref. 7). The simplest variant (and usually the one requiring the most space) is the *region octree*. In a region octree, a cubic domain is split at each node in eight equal sub-cubes. The decomposition continues recursively for each node that requires a further spatial refinement. Leaves of the octree represent empty or solid regions of space. The major drawbacks in the use of a region octree for representing the geometry are that it is an approximation (unless the object is constituted by mutually orthogonal planar faces only), and its size. Another octree variant is the *PM octree* (PM stands for Polygonal Map); each leaf corresponds to a single vertex, edge or face. The only exception is that a leaf may contain more than one edge (face) if all edges (faces) are incident on the same vertex. A PM octree usually requires much less storage than a region octree (Refs. 8 and 9), and it permits an exact representation of polyhedral models. PM octrees support a large suite of modeling operations, and several methods are known to convert CSG or Brep models to PM octrees and vice-versa (Ref. 7).

A *K-d-tree* (Ref. 10) (where K is the dimension of the domain space; in our case we will deal with 3-d-trees) is a binary tree in which internal nodes partition the space by a cut hyperplane defined by a value in one of the K dimensions, and external nodes, or *buckets*, store the points in the resulting hyper-rectangles of the partition. They allow $O(\log n)$ insertion, deletion and point query operation. The semi-dynamic variant introduced in Ref. 11 allows constant expected time deletion, undeletion, nearest-neighbor searching and fixed-radius near-neighbor searching. When used to model the geometry of an object, they suffer some of the same disadvantages of region octrees in that they can only approximate the geometry of general objects. However, they have been successfully used in precomputing viewpoint-independent visibility and illumination information for large models of buildings to speed up a successive interactive walk-through phase (Refs. 12-14). Changes in the environment would require a new processing of the visibility information.

Binary Space Partitioning Trees (BSPT) (Ref. 15) are a generalization of K -d-trees. In a BSPT, the cutting planes associated with the internal nodes are not constrained to be orthogonal. Constructing a BSPT representation of one or more polyhedral objects involves encoding the polygonal faces into a binary tree of cutting planes. Notice that affine transformations on the encoded object or groups of objects do not change the tree structure, but requires only the application of a transformation to the plane equations. The spatial ordering encoded in the tree can be exploited for the speedup of intersection and visibility computation. Moreover, when a BSPT is properly constructed it offers a *multi-resolution* representation of the object: as one descends a path in the tree, the bounded region decreases its size monotonically. The effectiveness of this technique in practice has been proven in several applications.

Unstructured tetrahedral grids are extensively used in finite element analysis and triangulations are pervasive in computational geometry. The use of simplicial complexes as a general approach to representing geometric shape has been advocated by several authors (Refs. 16 and 17). We are currently investigating the use of three-dimensional *regular* (a *weighted* variant of Delaunay) *triangulations*, coupled with efficient point location data structures, to provide the functionalities needed in a synthetic environment. A representation based on Hierarchical Simplicial Complexes has recently been proposed (Refs. 18 and 19) (see also an extension to a hierarchy of cell complexes (Ref. 20)). In our scheme, the top-level triangulation subdivides the space in tetrahedral cells (with the associated search structure). Each cell contains an object (or a group of objects), and these are described by other triangulations. The scheme can be recursively used to achieve the level of detail in the decomposition needed by the application. Using a hierarchy of triangulations has several advantages. Regular triangulations are acyclic with respect to visibility ordering (Ref. 21). This means that, given a viewpoint, it is possible to order the cells of the complex in a sequence such that if A obstructs B, then B precedes A in the sequence. This property is preserved in a hierarchy defined as above. Moreover, a hierarchy of simplicial complexes naturally defines a multi-resolution representation of the geometry, and allows fast point location and other types of queries when associated with appropriate data structures.

A three-dimensional regular triangulation can be built incrementally via point insertion and topological flipping. Additionally, a history DAG can be used to allow efficient point location and associated queries. Moreover, a *Power* (weighted Voronoi) *diagram* (eventually of order k), can be easily computed to allow fast answers to closest point queries. We are considering the use of several variations of this approach. A possible version of the data structure consists of using a Delaunay triangulation that conforms to objects boundaries. When a new object is inserted in the environment, the triangulation is updated to accommodate all its faces. Methods to construct conforming Delaunay triangulations in 2D and 3D by inserting extra points are known (Refs. 22-26). However, while a polynomial bound on the number of extra points is known for the two-dimensional case ($O(m^2n)$, for m edges and n vertices (Ref. 23)), whether such a bound exists for the three-dimensional case is still an open problem. Weatherhill (Ref. 26) reports statistics from experiments with the three-dimensional

algorithm that show a reasonable behavior on models used in real applications. This scheme is restricted to polyhedral objects (curved surface objects can be approximated by small planar faces).

In a different scheme, the triangulation is used as the domain for implicit piecewise-algebraic surfaces (Refs. 27-31). In each tetrahedron containing a piece of some object's boundary, a polynomial function $f(x,y,z)$ of degree n (where n is usually small, often $n \leq 3$) is defined, and the piece of surface is implicitly given by $f(x,y,z) = 0$. The surface patches can be made to join with some degree of continuity (for example, C^1 or C^2), by using interpolants of appropriate degree. For $n = 1$, in particular, the surface pieces are planar polygons, and the representation resembles the PM octree, where the cubic cells are replaced by tetrahedral ones. This approach has the advantage of allowing a compact representation for curved objects. Moreover, it makes possible the use of a more "relaxed" condition on the conforming triangulation. This needs not to conform to the object boundary, but simply contain a set of tetrahedra suitable to define the piecewise surface. An implicit piecewise-algebraic surface is suitable for interactive local control. In a three-dimensional synthetic environment, one could very naturally make use of three-dimensional widgets to allow an intuitive control on the surface shape.

Networked Distribution and Collaboration

We advocate the approach of integrating a collection of function-specific tools into a distributed and extensible environment where tools can easily use facilities provided by other tools (Refs. 34-36).

Isolation of functionality makes the environment modular and makes tools easy to develop and maintain. Distribution lets us benefit from the cumulative computation power of workstation clusters. Tool-level cooperation allows us to exploit the commonality that is inherent to many scientific manipulation systems. An enabling infrastructure of communication and interaction tools, display and visualization facilities, symbolic processing substrates, and simulation and animation tools saves avoidable re-implementation of existing functionality, and speeds up the application development.

The Shastra environment consists of multiple interacting tools. Some tools implement scientific design and manipulation functionality (the Shastra Toolkits). Other tools are responsible for managing the collaborative environment (Kernels and Session Managers). Yet others offer specific services for communication and animation (Service Applications). Tools register with the environment at startup, providing information about the kind of services that they offer (Directory), and how and where they can be contacted for those services (Location). The environment supports mechanisms to create remote instances of applications and to connect to them in client-server or peer-peer mode (Distribution). In addition, it provides facilities for different types of multi-user interaction ranging from master-slave blackboarding (Turn Taking) to synchronous multiple-user interaction (Collaboration). It implements functionality for starting and terminating collaborative sessions and for joining or leaving them. It also supports dynamic messaging between different tools. Tools are thus built on top of the abstract Shastra layer which is depicted in Fig. 1.

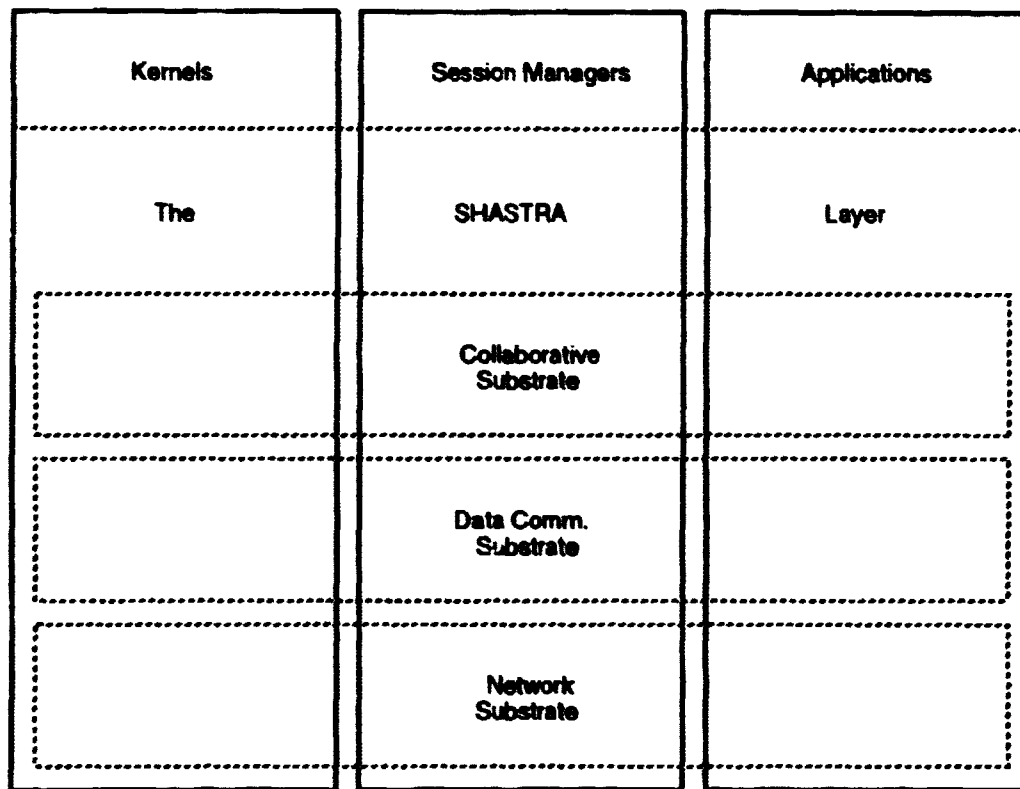


Figure 1 - The Shastra Layer: A connection, communication and collaboration management substrate. Shastra tools inter-operate using facilities provided by this layer.

The SHAstra collaborative scientific environment provides mechanisms to support a variety of multi-user interactions spanning the range from demonstrations and walk-throughs to synchronous multi-user collaboration. In addition, it facilitates synchronous and asynchronous exchange of multimedia information which is useful to successfully communicate at the time of design, and to share the results of scientific tasks, and often necessary to actually solve problems. The infrastructure provides facilities to distribute the input of low computation tasks - to obtain the parallelism benefit of distribution, and the output of compute intensive tasks - to emphasize sharing of resources among applications. It provides a convenient abstraction to the application developer, shielding him from lower level details, while providing him with a rich substrate of high level mechanisms to tackle progressively larger problems.

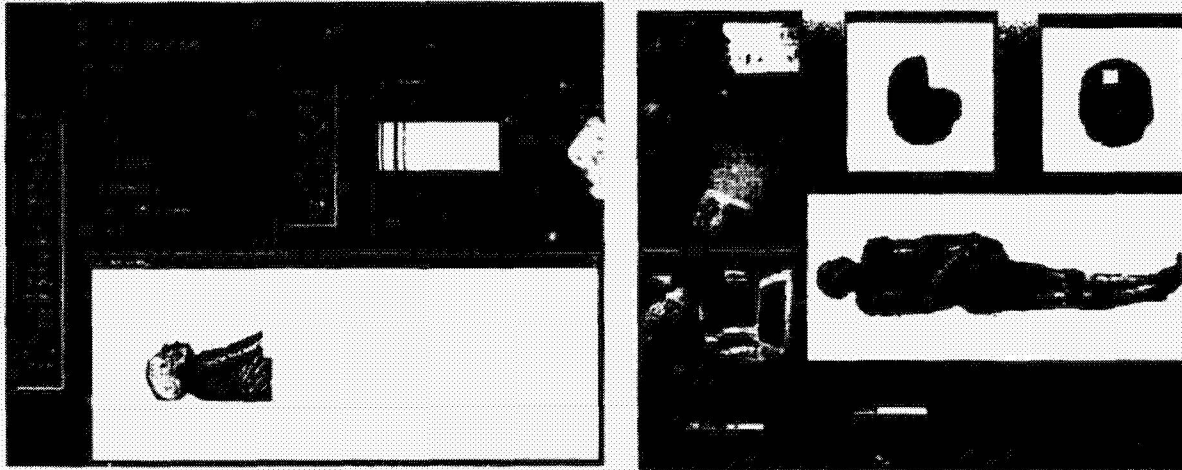


Figure 2 - Distributed volume visualization from a large data set (left). Shared Visualization (right). A group of researchers uses Sha-Poly to share volume visualization images of a head with cutaways (top center and right) and a cadaver (center).

A teleconferencing approach to modeling and analysis of empirical data is presented in Ref. 35, where the authors hypothesize about a collaborative scientific visualization environment. A discussion of an interactive visualization environment for three-dimensional imaging is presented in Ref. 36, where the authors adapt the electron microscope to perform as a computer peripheral. An environment like Shastra makes it convenient to build collaborative visualization and manipulation facilities that support resource sharing in a distributed setting.



Figure 3 - Collaborative customized hip implant design (left). A designer uses the Shastra collaborative tool Shilp to interactively create a geometric model of a hip implant (top right) by generating cross-sectional contours of the implant (bottom center and right) from a sectional model of the femur (center) created in Vaidak (another toolkit of the Shastra suite). A video conference supports communication (right). Several steps in the design optimization.

Shastra Architecture

Tools in Shastra are built with the underlying idea of inter-tool cooperation. Every tool is abstractly composed of three layers. The Core is accessed through any of the Interfaces via a Mapper.

The application-specific Core implements the functionality offered by the tool. Above the Core is a functional Interface Mapper that invokes functionality embedded in the Core in response to requests from the Graphical User Interface, ASCII Interface or the Network Interface. It also maps requests to alter the user interface or to send messages on the Network Interface. The Mapper is essentially a command interpreter that invokes registered event handlers when events of interest occur. Tools register event handlers with the Mapper for events they are interested in and unregister those that cease to be of interest. The separation of Core and Interface, that of function and interface, makes it easy to build multi-user systems since it enables the maintenance and display of shared state at a user interface *via* remote commands in a distributed system.

The GUI is application-specific. The ASCII interface is a shell-like front end for the tool. Tools communicate with other tools in the environment *via* the Shastra substrate, through an abstract Network Interface. This implements the underlying messaging system that provides connection and transport facilities. The Network Interface multiplexes multiple simultaneous network connections and implements the different application level communication protocols. Functionality available at a network interface is relayed to the communication substrate using a signature that specifies callback functions for the different kinds of network events in which the tool is interested. The signature provides an abstract interface to remote systems and describes functionality offered by the tool. It also serves as a regulatory mechanism since different levels of service can be offered at different interfaces by specifying the appropriate signatures.

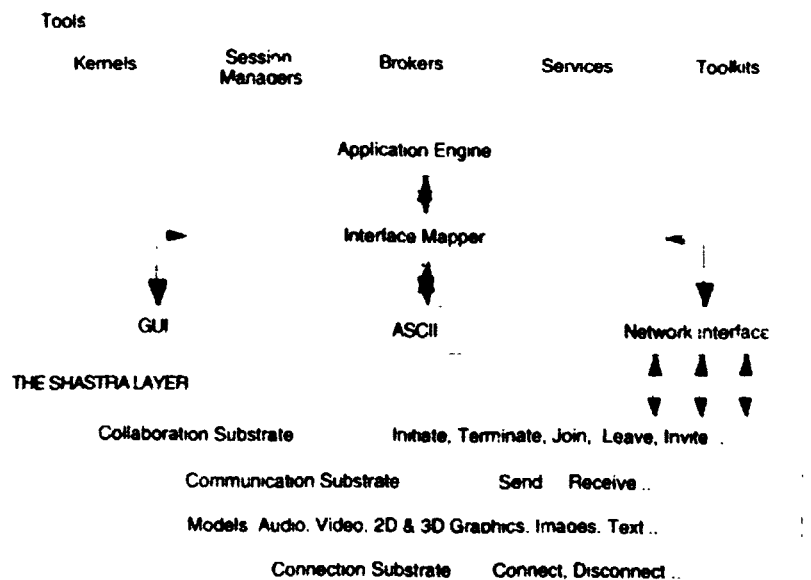


Figure 4 - High level architecture of a tool in the Shastra environment.

To take advantage of the integration facilities of the infrastructure, the Core uses the Network Interface to access functionality already implemented in other tools. The main benefit from this setup is modularity and reuse - tools isolate the functionality they offer and provide a functional interface to peers. The high level block architecture of tools in Shastra is depicted in Fig. 4. The architecture makes it easy for tools to connect to other tools and request operations, synchronously as well as asynchronously.

These architectural guidelines accord us the benefit of uniformity since all tools are built upon a common infrastructure and have identical connection, communication and collaboration mechanisms.

The concept of cooperation awareness thus pervades the architecture. The entire set of connected Network Interfaces of Shastra tools manifests itself as the abstract Shastra layer at runtime (see Fig. 4). It maintains the collaborative environment, provides access to functionality of different systems, and provides facilities for initiating, terminating, joining, leaving and conducting collaborations. The connected network interfaces of Shastra tools comprise a distributed virtual machine on which we build problem solving applications.

The enabling substrates use the event paradigm to provide functionality. Tools use the application programming interface of the substrate to cause request messages to be sent over connections. Tools interested in any event register handler functions for it with the Mapper. The handler functions are invoked when that event is received. This allows tools to take action appropriate to the event when it occurs.

Session Model

In this model, a Session is the unit of collaborative activity. A Session is essentially a Context without an Interface. Session Model based collaborative tools are implemented in our Collaboration Model by instantiating a Session which causes the setting up of connected shared Contexts in multiple tools. These shared Contexts are collaborative task-aware. Events that are associated with low computation tasks are routed to the Session Context which relays them to all shared Contexts. Events that are associated with compute-intensive tasks are acted upon in the tool Context, and the associated Triggers are routed to the Session Context. Context State changes generate Triggers that are routed to tools and update views at their Interfaces. The implementation is depicted in Fig. 5.

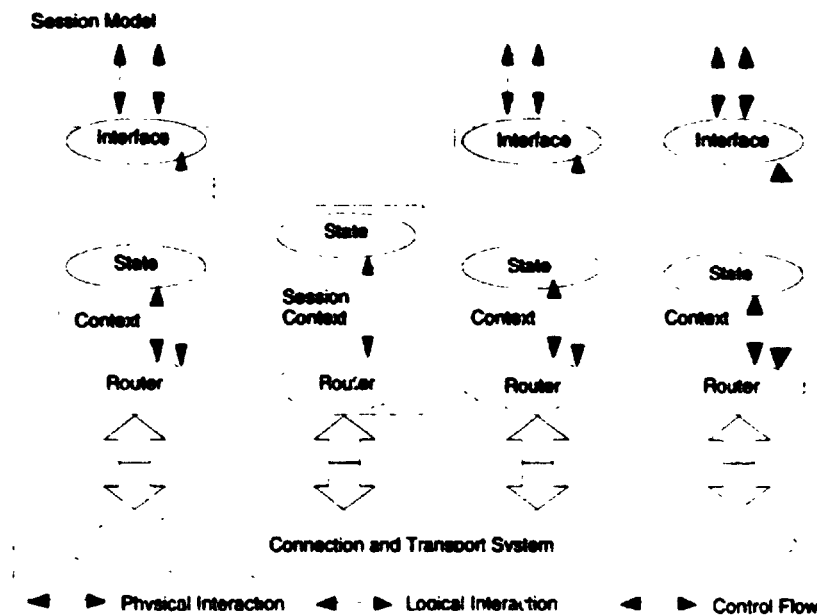


Figure 5 - Session model of collaboration.

Since Sessions are collaborative task-aware, they can choose between centralized and replicated data management facilities based on the number of sites in the collaboration, degree of dependence between collaborative tasks, and performance of the underlying mechanisms. Collaborative tasks are thus implemented in the shared Context and State of a Session. Simultaneous interaction is supported from multiple distributed interfaces.

A major advantage of this approach is that Sessions can be made persistent since they are delinked from user level tools and interfaces. They can be saved and restarted and thus support asynchronous and synchronous collaborative interaction. Also, participating in collaborative tasks is further simplified since tools do not have to keep track of group membership or set up routing information. Tools create Contexts that are shared with the Session Context when they join a Session and tear them down when they leave.

Runtime Environment

Collaborative Sessions, or Sessions, are instances of synchronous multi-user collaborations or conferences in the Shastra environment. A collaboration in Shastra consists of a group of cooperating tools regulated by a Session Manager, the conference management tool of Shastra. One Session Manager runs per collaborative session. It maintains the session and handles details of connection and session management, interaction control and access regulation. It keeps track of membership of the collaborative group and serves as a repository of the shared objects in the collaboration. It supports a multicast facility needed for information exchange in a synchronous multi-user conferencing scenario. It has a constraint management subsystem that resolves conflicts that arise as a result of multi-user interaction, enabling maintenance of mutual consistency of operations. It has a regulatory subsystem that controls synchronous multi-party interaction and provides a floor control facility based on turn-taking. Every Session Manager implements functionality to service the following session control requests:

- Invite - Request to invite a tool to an ongoing session.
- Join - Request to join an ongoing session.
- Remove - Request to remove a tool from a session.
- Leave - Request to leave a session of which the tool is a member.
- End - Request to terminate a collaborative session.

It also serves the following interaction control requests:

- Format - Request to set session format.
- Capabilities - Request to set access regulation capabilities.
- Interaction Mode - Request to set interaction mode for the session.
- Request Floor - Request to get floor control for the session.
- Release Floor - Request to release floor control for the session.
- Assign Floor - Request to assign floor control for the session.

A collaborative session in Shastra is started by a tool when it sends the Session Start message to the local Kernel. This causes the instantiation of a Session Manager for the incipient session. The initiating tool becomes the Session Leader. A tool sets the session format using the Format message. Sessions may be Formal, where participation is by invitation only, or Informal where any tool can dynamically join the conference. The Leader assigns capabilities of other participants for collaborative activity in the session using the Capabilities message.

The interaction mode for a session is specified using the Interaction Mode message. Interaction can occur in the Regulated or Free mode. In the Regulated mode, tools request and relinquish the floor using Request Floor and Release Floor messages. The leader can explicitly assign floor control using the Assign Floor message. In the Free mode, interaction is regulated *via* capabilities assigned to session participants. Capabilities are described in a later section. Other tools are invited to participate in a session by sending them the Invite request *via* the Kernel. Tools can dynamically join ongoing sessions by sending the Join message to the relevant Session Manager *via* the Kernel. The Session Manager uses session format information to control dynamic incorporation of tools. The Leader can

remove a participating tool from the session using the Remove message. Tools can discontinue participation in the session by sending the Leave message to the Session Manager. A session is terminated by the Leader using the End message.

Application-specific Session Managers for different collaborative tasks are created from the basic Session Manager that provides application independent connection, communication and collaboration control facilities. Such session managers support additional messages for collaborative operations specific to the application.

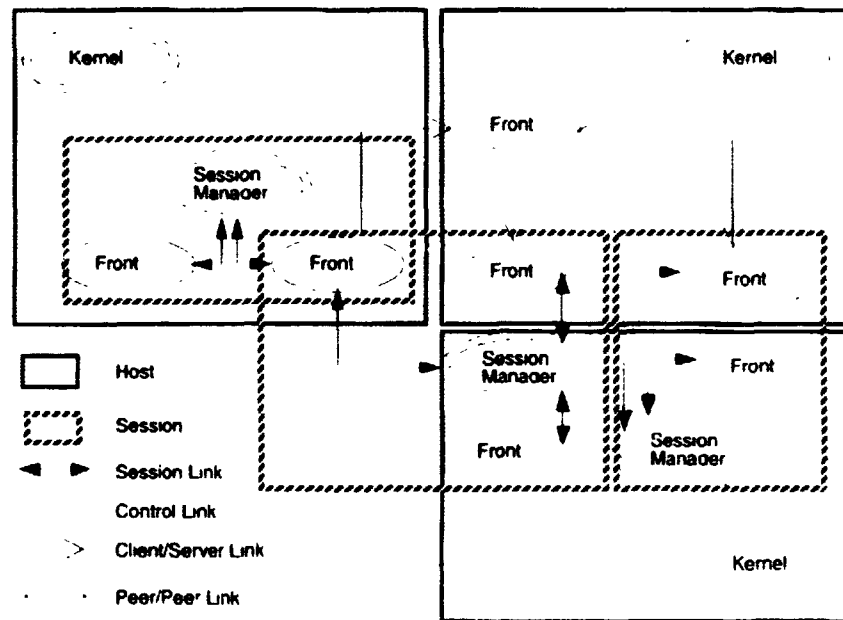


Figure 6 - Information flow in the Shastra environment.

References

1. MacDonald, L. and Vince, J., *Interacting with Virtual Environments*, Wiley Professional Computing, Wiley, 1994.
2. Thalmann, N. M. and Thalmann, D., *Virtual Worlds and Multimedia*, Wiley Professional Computing, Wiley, 1993.
3. Wexelblat, A., *Virtual Reality: Applications and Explorations*, Academic Press Professional, 1993.
4. Special Issue on Virtual Environments, *IEEE Computer*, Vol. 28, No. 7, July 1995.
5. Special Issue on Virtual Reality, *IEEE Computer Graphics and Applications*, Vol. 14, No. 1, Jan. 1994.
6. Special Issue on Virtual Reality, *IEEE Spectrum*, Vol. 30, No. 10, Oct. 1993.

7. Samet, H., *The Design and Analysis of Spatial Data Structures*, Addison-Wesley, Reading, MA, 1990.
8. Ayala, D., Brunet, P., Juan, R. and Navazo, I., "Object Representation by Means of Nominal Division Quadrees," *ACM Trans. Graph.*, Vol. 4, No. 1, 1985, pp. 41-59.
9. Brunet, P., and Navazo, I., "Solid Representations Using Extended Octrees," *ACM Trans. Graph.*, Vol. 9, 1990, pp. 170-197.
10. Bentley, J. L., "Multidimensional Binary Search Trees Used for Associative Searching," *Communications of the ACM*, Vol. 18, No. 9, 1975, pp. 509-517.
11. Bentley, J. L., "K-d Trees for Semi-dynamic Point Sets," *Proc 6th Annual ACM Symposium on Comput. Geom.*, 1990, pp. 187-197.
12. Teller, S. J. and Sèquin, C. H., "Visibility Preprocessing for Interactive Walkthroughs," *Comput. Graph.*, Vol. 25, No. 4, July 1991, pp. 61-69.
13. Teller, S. and Hanrahan, P., "Global Visibility Algorithms for Illumination Computations," *Computer Graphics Proc., Annual Conference Series, Proc. SIGGRAPH '93*, ACM SIGGRAPH, 1993, pp. 239-246.
14. Funkhouser, T. A. and Sèquin, C. H., "Adaptive Display Algorithm for Interactive Frame Rates During Visualization of Complex Virtual Environments," *Computer Graphics Proc., Annual Conference Series, Proc. SIGGRAPH '93*, ACM SIGGRAPH, 1993, pp. 247-254.
15. Fuchs, H., Kedem, Z. M. and Naylor, B., "On Visible Surface Generation by a priori Tree Structures," *Comput. Graph.*, Vol. 14, No. 3, 1980, pp. 124-133.
16. Paoluzzi, A., Bernardini, F., Cattani, C. and Ferrucci, V., "Dimension-Independent Modeling with Simplicial Complexes," *ACM Transactions on Graphics*, Vol. 12, No. 1, Jan. 1993, pp. 56-102.
17. Edelsbrunner, H., "Modeling with Simplicial Complexes," *SIGGRAPH 94 Course Notes*, ACM SIGGRAPH, July 1994.
18. De Floriani, L., Falcidieno, B. and Pienovi, C., "Structured Graph Representation of a Hierarchical Triangulation," *Comput. Vision Graph. Image Process*, Vol. 45, No. 2, Feb. 1989, pp. 215-226.
19. De Floriani, L., "A Pyramidal Data Structure for Triangle-Based Surface Representation," *IEEE Computer Graphics and Applications*, Vol. 9, March 1989, pp. 67-78.
20. Bertolotto, M., Bruzzone, E. and De Floriani, L., "Acyclic Hierarchical Simplicial Complexes," *Proc 5th Canadian Conference on Comput. Geom.*, 1993, Waterloo, Canada, pp. 279-284.
21. Edelsbrunner, E., "An Acyclicity Theorem for Cell Complexes in d Dimensions," *Proc 5th Annual ACM Symposium on Comput. Geom.*, 1989, pp. 145-151.
22. Weatherill, N. P., "Integrity of Geometrical Boundaries in the Two-Dimensional Delaunay Triangulation," *Communications in Applied Numerical Methods*, Vol. 6, No. 2, Feb. 1990, pp. 101-109.
23. Edelsbrunner, H. and Tan, T. S., "An Upper Bound for Conforming Delaunay Triangulations," *Proc 8th Annual ACM Symposium on Comput. Geom.*, 1992, pp. 53-62.

24. Hansen, A. J. and Levin, P. L., "On Conforming Delaunay Mesh Generation," *Advanced Engineering Software*, Vol. 14, No. 2, 1992, pp. 129-135.
25. Tan, T. S., "An Optimal Bound for Quality Conforming Triangulations," *Proc 10th Annual ACM Symposium on Comput. Geom.*, 1994, pp. 240-249.
26. Weatherill, N. P., "Efficient Three-Dimensional Delaunay Triangulation with Automatic Point Creation and Imposed Boundary Constraints," *International Journal for Numerical Methods in Engineering*, Vol. 37, June 1994, pp. 2005-2039.
27. Bajaj, C., "Surface Fitting with Implicit Algebraic Surface Patches," *Topics in Surface Modeling*, H. Hagen (ed), SIAM Publications, pp. 23-52.
28. Bajaj, C., "The Emergence of Algebraic Curves and Surfaces in Geometric Design," *Directions in Geometric Computing*, R. Martin (ed), Information Geometers Press, Winchester, UK, 1993, pp. 1-29.
29. Bajaj, C., Bernardini, F. and Xu, G., *Adaptive Reconstruction of Surfaces and Scalar Fields From Dense Scattered Trivariate Data*, Computer Science Technical Report CSD-TR-95-028, Purdue University, 1995.
30. Bajaj, C., Bernardini, F. and Xu, G., "Automatic Reconstruction of Surfaces and Scalar Fields From 3D Scans," *Computer Graphics Proc, Annual Conference Series, Proceedings of SIGGRAPH 95*, ACM SIGGRAPH, 1995, pp. 109-118.
31. Bajaj, C., Chen, J. and Xu, G., "Modeling with Cubic A-patches," *ACM Transactions on Graphics* (to appear).
32. Anupam, V. and Bajaj, C., "Collaborative Multimedia in Scientific Design," *IEEE Multimedia Journal*, Vol. 1, No. 2, 1994, pp. 39-49.
33. Anupam, V. and Bajaj, C., "Shastra: An Architecture for Development of Collaborative Applications," *International Journal of Intelligent and Cooperative Information Systems*, Vol. 3, No. 2, 1994, pp. 155-166.
34. Anupam V., Bajaj, C. and Schikore, D., "Distributed and Collaborative Volume Visualization," *IEEE Computer*, Vol. 27, No. 7, 1994, pp. 37-43.
35. Carlbom, I., Hsu, W., Klinkner, G., Szeliski, R., Waters, K., Doyle, M., Gettys, J., Harris, K., Levergood, T., Palmer, R., Picart, M., Terzopoulos, D., Tonnesen, D., Vannier, M. and Wallace, G., "Modeling and Analysis of Empirical Data in Collaborative Environments," *Communications of the ACM*, Vol. 41, No. 6, 1992, pp. 73-84.
36. Mercurio, P., Elvine, T., Young, S., Cohen, P., Fall, K. and Ellisman, M., "The Distributed Laboratory," *Communications of the ACM*, Vol. 41, No. 6, 1992, pp. 54-63.

21037

Information Visualization - Beyond Traditional Engineering

James J. Thomas
Applied Physics Center
Battelle Pacific Northwest Laboratories
Richland, WA

**NEXT
DOCUMENT**

222
6251
1 2 1 1

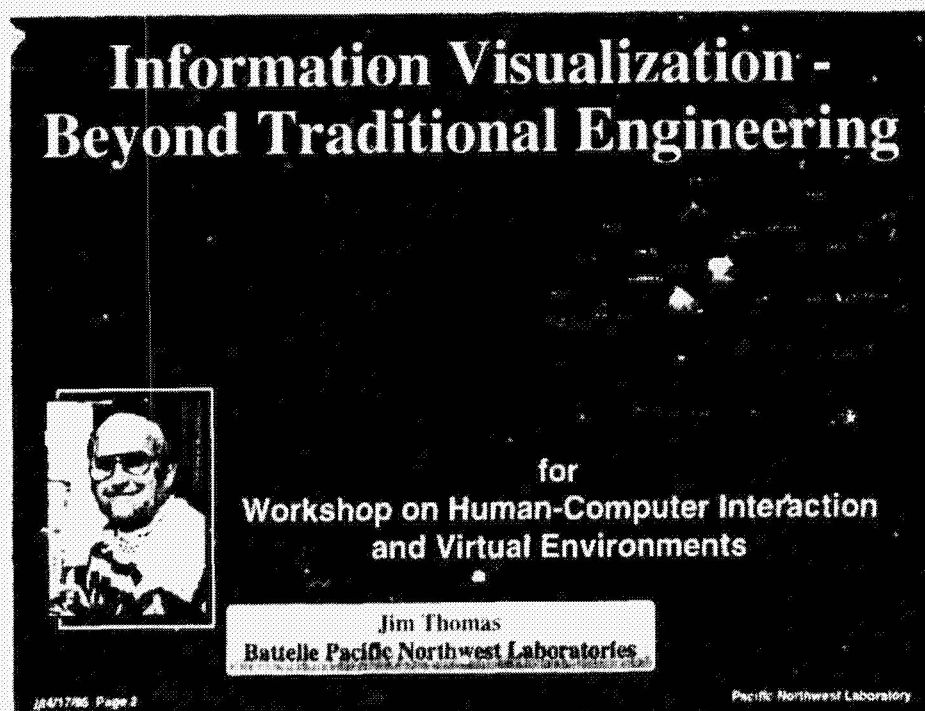
INFORMATION VISUALIZATION - BEYOND TRADITIONAL ENGINEERING

James J. Thomas
Technology Manager
Applied Physics Center
Battelle Pacific Northwest Laboratories
Richland, WA 99352

I would like to thank the Workshop committee for the opportunity to address this audience at this important workshop on Human-Computer Interaction and Virtual Environments.


So far you have heard a lot about the traditional devices and interactions within virtual environments. This presentation will address a different aspect of the human-computer interface; specifically the *human-information* interface.

This interface will be dominated by an emerging technology called Information Visualization. Information goes beyond the traditional views of computer graphics, CADs and enables new approaches for engineering.



**Information Visualization -
Beyond Traditional Engineering**

for
**Workshop on Human-Computer Interaction
and Virtual Environments**


Jim Thomas
Battelle Pacific Northwest Laboratories

JA471786 Page 2 Pacific Northwest Laboratory

TAKE HOME MESSAGE FOR HCI AND COMPLEX INFORMATION SPACES

The take home message is simple.

Information visualization is the visual and interaction technologies supporting analysis of all forms of information including text, images, diagrams, procedures, marketing materials and product quality information

Information visualization will add value by providing the holistic approach to engineering quality products.

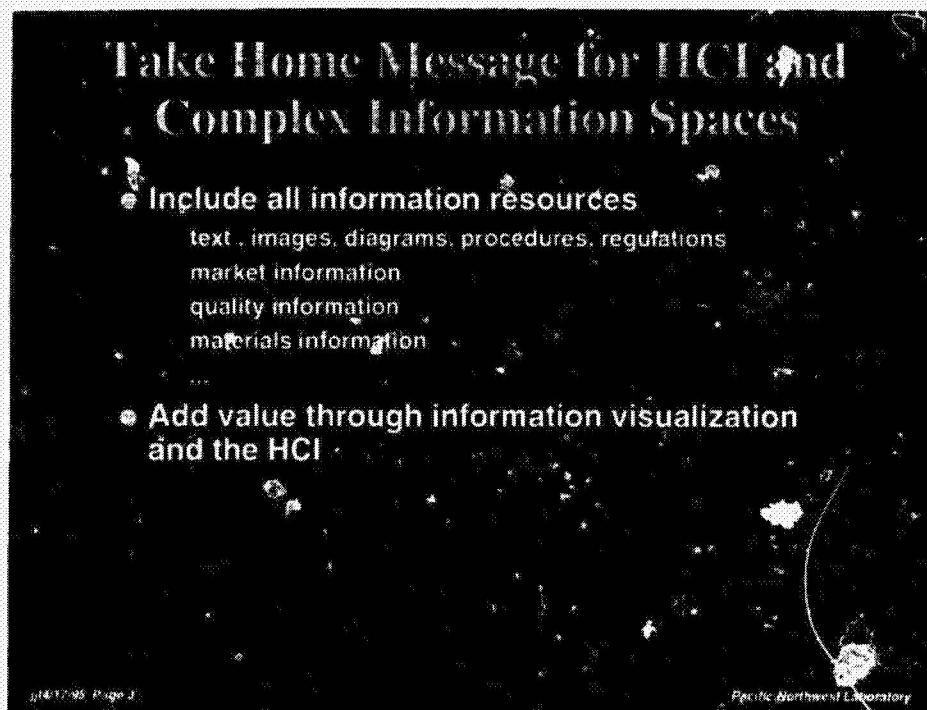


Figure 1

TAKE HOME MESSAGE FOR HCI AND COMPLEX INFORMATION SPACES

However, to enable this fundamental change in the engineering process, our R&D must focus on all forms of information spaces.

Government and industry alike must change their traditional views of scientific visualization. Information visualization will become an integral part of the entire product cycle.

Take Home Message for HCI and Complex Information Spaces

- **Include all information resources**
 - ◆ text, images, diagrams, procedures, regulations
 - ◆ market information
 - ◆ quality information
 - ◆ materials information
 - ◆ ...
- **Add value through information visualization and the HCI**

Focus R&D on total information spaces for today's and tomorrow's industry

PNL/STP Page 4 Pacific Northwest Laboratory

Figure 2

OUTLINE

I would like to organize this talk in two sections.

First, let's look at the progress to date over the last two decades.

Then, let's look at emerging technology for information visualization: specifically, technology to visualize masses of text. This is a fundamentally new approach and is right at the core of information visualization.

Before we start, I want to clarify what is meant by information visualization.

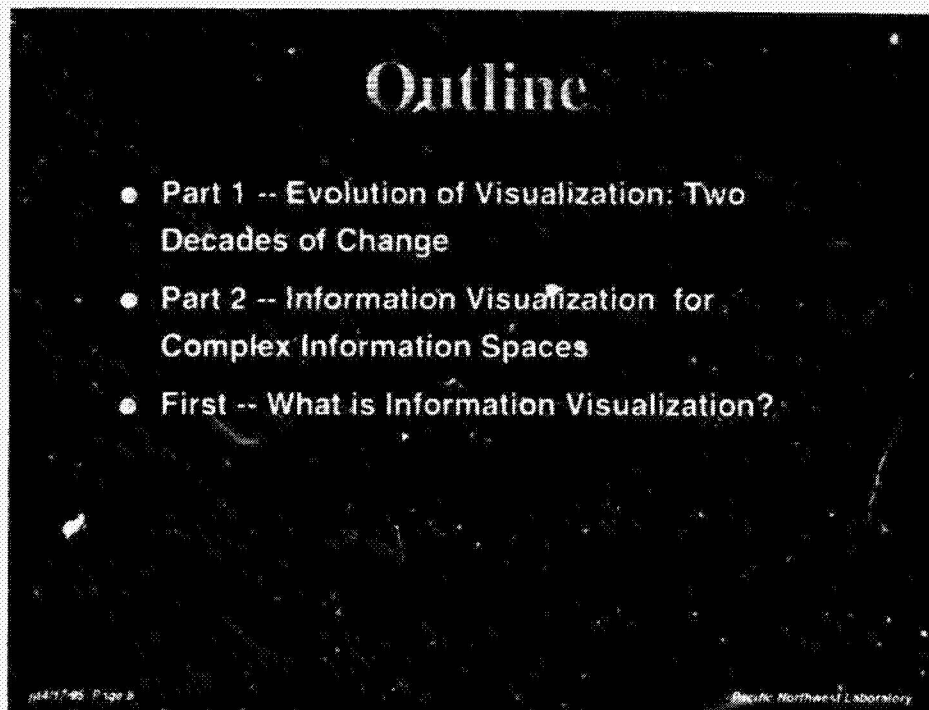


Figure 3

WHAT IS INFORMATION VISUALIZATION?

Information visualization (IV) is not just visualizing the numbers, engineering diagrams, etc. IV goes way beyond traditional scientific visualization based on physical properties. IV specifically must visualize text, documents, sound, images and video in such a way that the human can rapidly interact with and understand the content structure of information entities. These entities are not based on math, physics or chemistry, but rather on concepts; yes, often fuzzy, often incomplete, and often related to other concepts in many dimensions.

It is indeed a high-dimensional fuzzy set of information entities that we deal with every day of our lives.

What is Information Visualization?

Information Visualization is the interactive visual interface between humans and their information resources.

- **Information Visualization is not just**
 - ◆ display of numbers - text, documents, video, sound, imagery, ...
 - ◆ theory or experiment results - libraries, directories, information spaces, offices, ...
 - ◆ local data -- networks and worlds of information
 - ◆ in the office - in the home, auto, air, ... (TV: through set boxes)
 - ◆ science - business, home, education, ...

4/8/17/95 Page 6 Pacific Northwest Laboratory

Figure 4

WHAT IS INFORMATION VISUALIZATION?

My colleague, Jim Wise, a cognitive psychologist, says that a primary goal of IV is to provide the presentation and interactions that match our trained perceptual capabilities. These capabilities are developed in our childhood and perfected as we mature both in the home and office settings. These capabilities also allow us to visualize and understand masses of information simultaneously. Look at our desks as only a part of the total information space in the office.

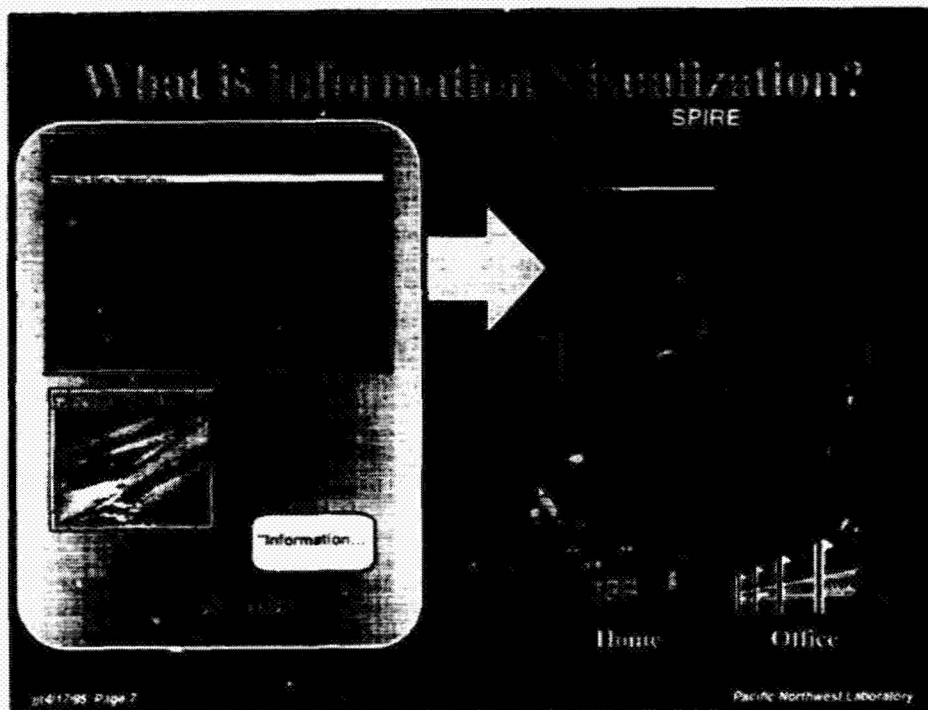


Figure 5

PART 1 - TWO DECADES OF CHANGE

Now let's look at what has happened in the last two decades. I think that you will be as surprised as I was when I spent some time reviewing our progress.

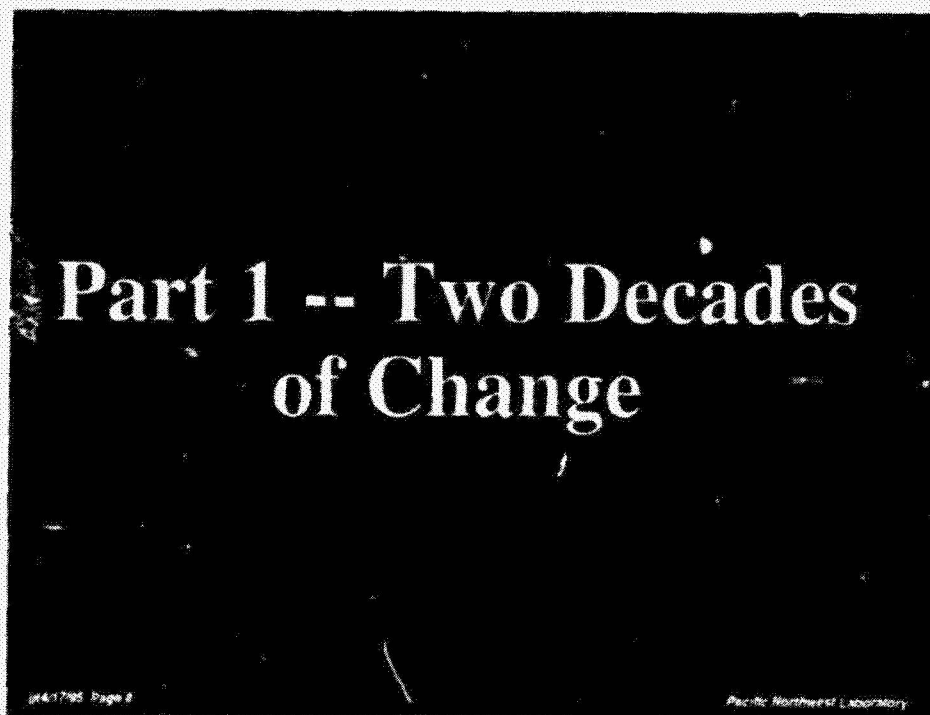


Figure 6

We are being driven by a dramatic change in our society. This was discussed by Alvin and Heidi Toffler in their book called War and Anti-War. If you have not read the book, I highly recommend it. It will enable you to think in a different perspective.

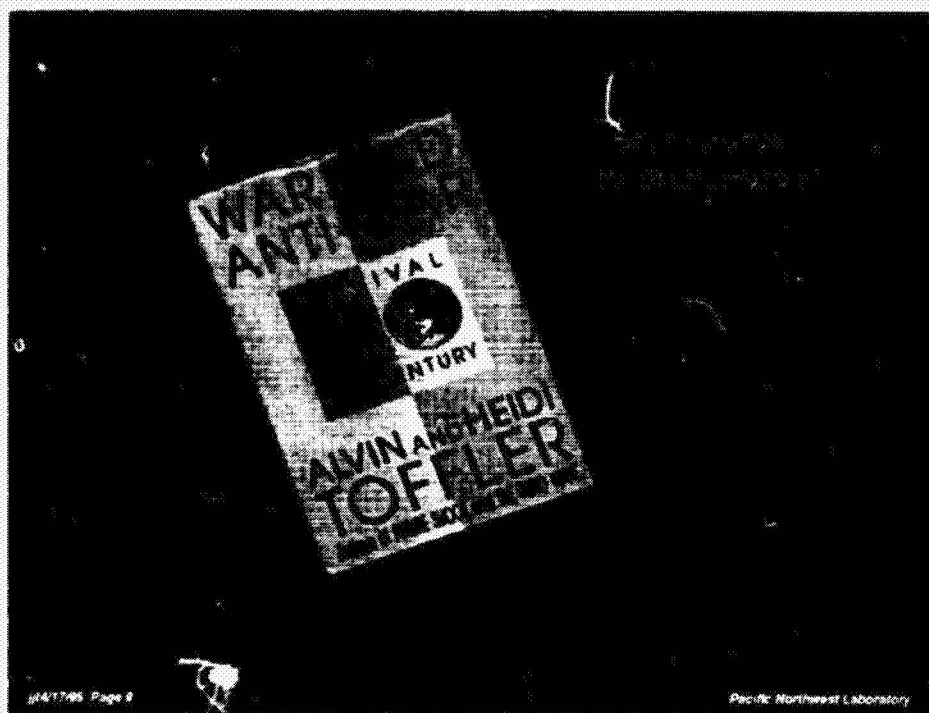


Figure 7

Leaving the geo-centric era behind, we are indeed entering the geo-information era.

The ability to deal with masses of information will be a key part of competitiveness within and between our societies.

Talk about quotes.

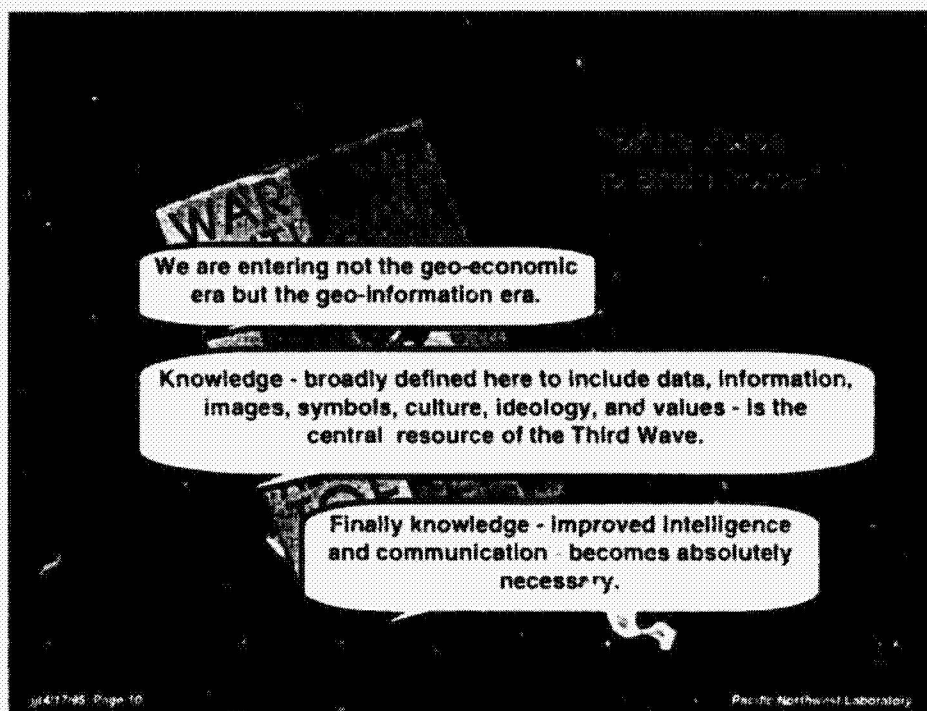


Figure 8

TWO DECADES OF CHANGE

Now that we have seen the motivations, where are we?

In the 1970's we invented scientific visualization that was driven by initial needs for CAD for the automobile and aircraft industries. Some of you in the audience are pioneers in this arena. Then technology was driven by the needs and \$\$\$ within the entertainment industry. This was followed by a continuing push to enable the use of multimedia sources of information. However, the primary source of information, TEXT, was not addressed. Until we address this key form of information we will see little change. Some may not agree.

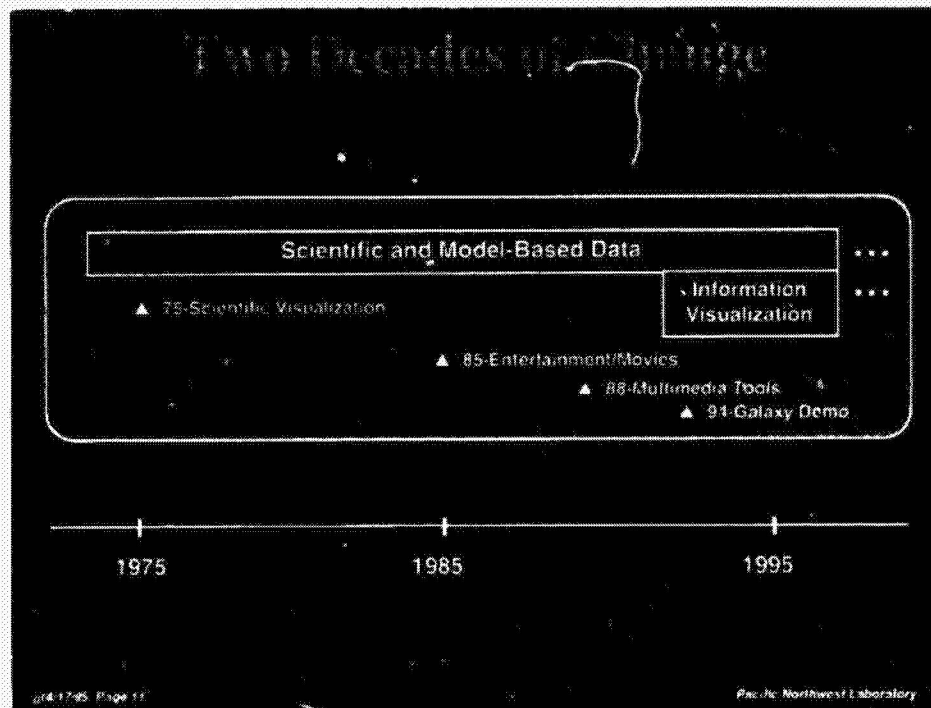


Figure 9

VISUALIZATION TECHNOLOGY EVOLUTION

Let's look at the changes in some visualization technology from the early 1980's through 1995. These are segments of film from contributors at the SIGGRAPH Film and Video show held each year. They are available and referenced from the Video Review below. These pieces are major contributions, and credits to the authors are provided.

Please note that the initial sequences took "CRAY" hours per frame to generate. Today these segments are almost always computed in real time; some are. Surprisingly, we see little change except for quality and speed.

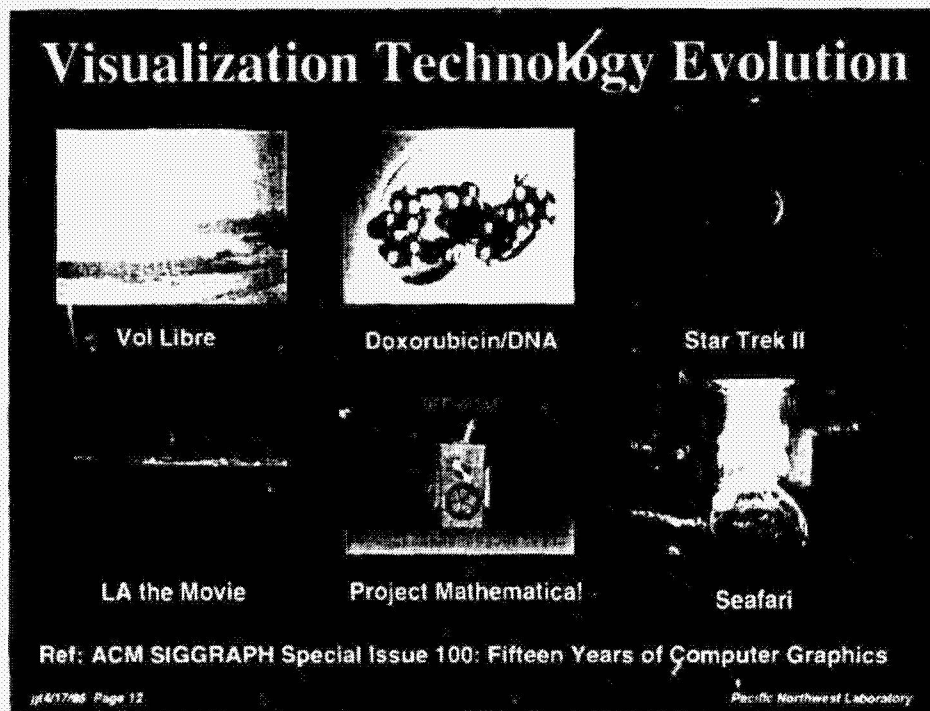


Figure 10

LAST TWENTY YEARS - OBSERVATIONS

You have already seen current technology in engineering visualization. What can we conclude?

- Faster and cheaper by two orders of magnitude - quality is definitely increased.
- It is being used effectively by scientists and engineers.
- The entertainment community has been the dominate driver.
- Many applications are being driven by the technology versus the application needs.
- We can deal with much larger volumes of data.

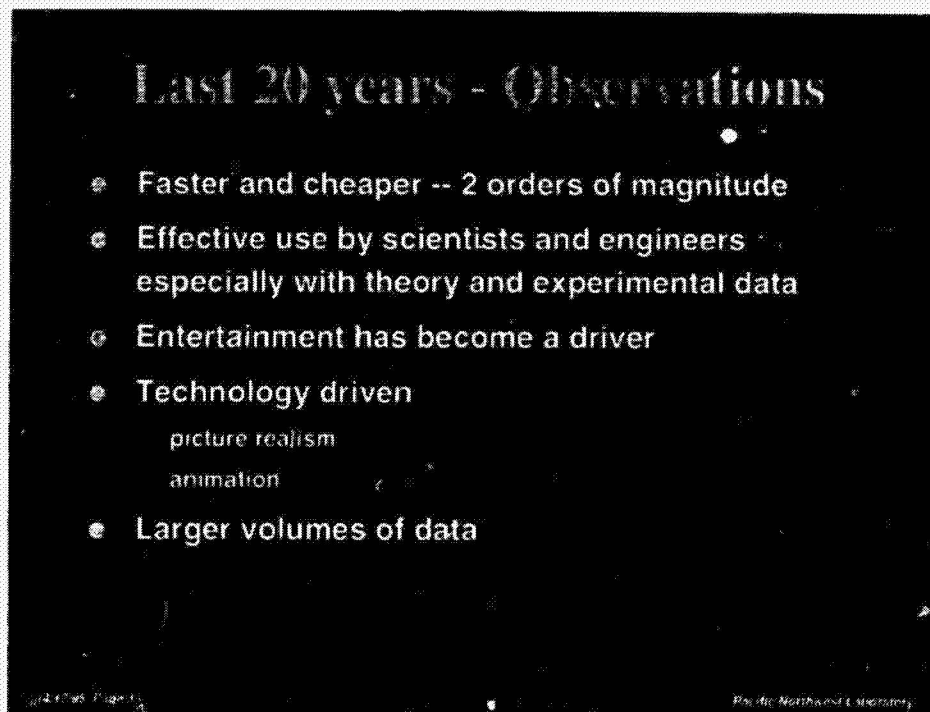


Figure 11

LAST TWENTY YEARS - OBSERVATIONS

- Interaction technologies are essentially the same with the possible exception of some of the VR interaction devices. It is windows; point and click, and WIMP.
- Today's visualization is largely based on the physical properties of materials and designs.
- There is little change in how we deal with information analysis. We can even simulate the old methods somewhat faster on the computer. With the changing data volumes, it is hard to determine any real progress.
- We mostly analyze information of known and complete structure.
- There is little-to-no technology for the inclusion of the total information spaces for analysis.

Last 20 years - Observations

- Interaction technologies haven't changed
- Visualization focused on theory and physical information
 - ◆ looking for exact solutions in fuzzy world
 - ◆ very little for the educator and policy maker
- Little change in how we analyze information
- Mostly analyze data of known structure

No application to improving understanding of total information resources

ptl417-95 Page 14 Pacific Northwest Laboratory

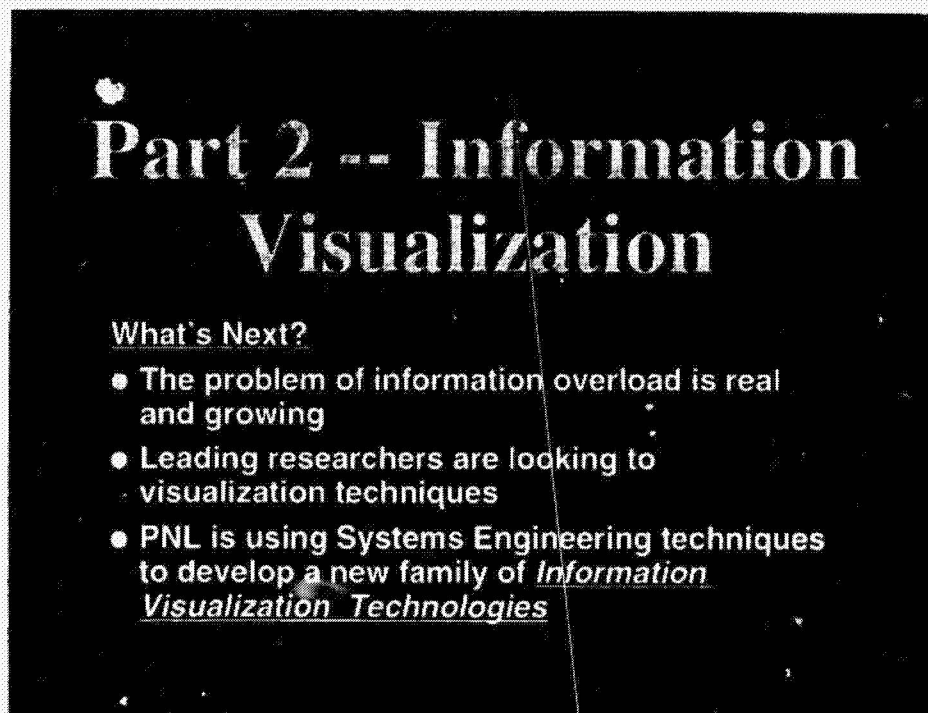
Figure 12

PART 2 - INFORMATION VISUALIZATION

Information overload is about to happen. You may say that you are swamped today. Consider what will happen when you have 100-1000 times the information available to you on any topic of your interest. That is the issue to be addressed by IV.

There are some leading researchers in the field attempting to address this.

PNL's approach is to conduct research while delivering today's technology, getting direct feedback from information analysts using a systems engineering methodology. This is the subject of another talk. This talk will focus on the resulting technology.



Part 2 -- Information Visualization

What's Next?

- The problem of information overload is real and growing
- Leading researchers are looking to visualization techniques
- PNL is using Systems Engineering techniques to develop a new family of Information Visualization Technologies

Figure 13

INFORMATION OVERLOAD

Is this your office?

Play video - illustrating a typical information analyst.

Information Overload

Key Factors

- 90% of all data in the world is in text form
- Implementation of World-Wide-Web, NII, and HPCC
- Growth of openly available information
- Need for real-time information analysis
- Increasing growth of business information access




Figure 14

PURPOSE OF INFORMATION VISUALIZATION

The purpose of IV is to:

Provide an enhanced method of analysis that enables discovery, understanding and presentation through the use of computer graphics and the interactive interface between the human and their information resources.

Note: Discovery, understanding and presentation are the three fundamental activities of an information analyst.

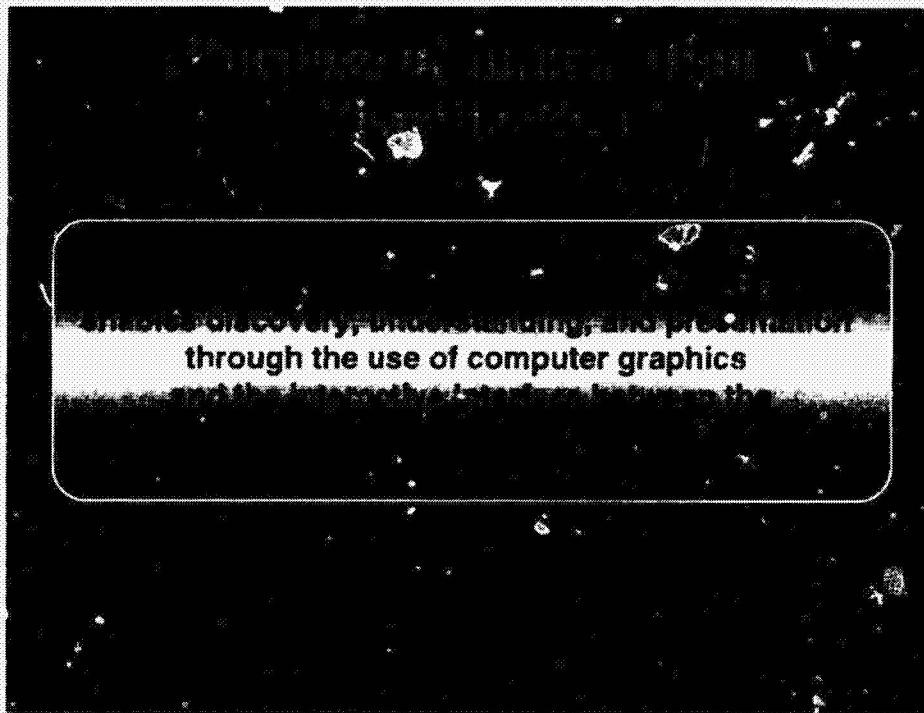


Figure 15

RESEARCH IN INFORMATION VISUALIZATION

Examples of significant contributions are from:

ALTA Analytics - Netmap. Their software provides that each phrase is related to other phrases through a vector map, which is illustrated. This works well for small numbers of single-dimensional information spaces. They have some nice interaction tools that help with the more complex information spaces.

Xerox Parc's Cone Trees provides an excellent approach for IV of hierarchical information.

Ben Schneiderman's scatter plots, demonstrated by the FilmFinder, is an excellent example of interacting with medium-sized data sets. This provided many clues leading to the following work from PNL.

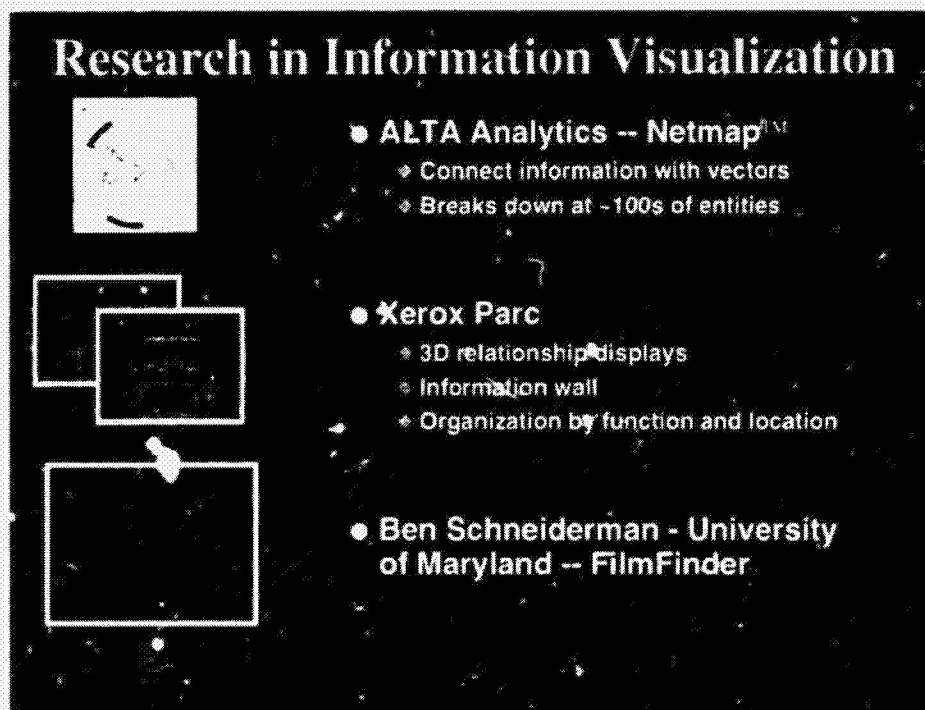


Figure 16

KEY TECHNOLOGIES FOR INFORMATION VISUALIZATION

There are at least ten fundamental technologies that support work in IV. These have been documented within a state-of-the-art report that is available if you are interested. Please contact me via e-mail (JJ_Thomas@pnl.gov) for a copy.

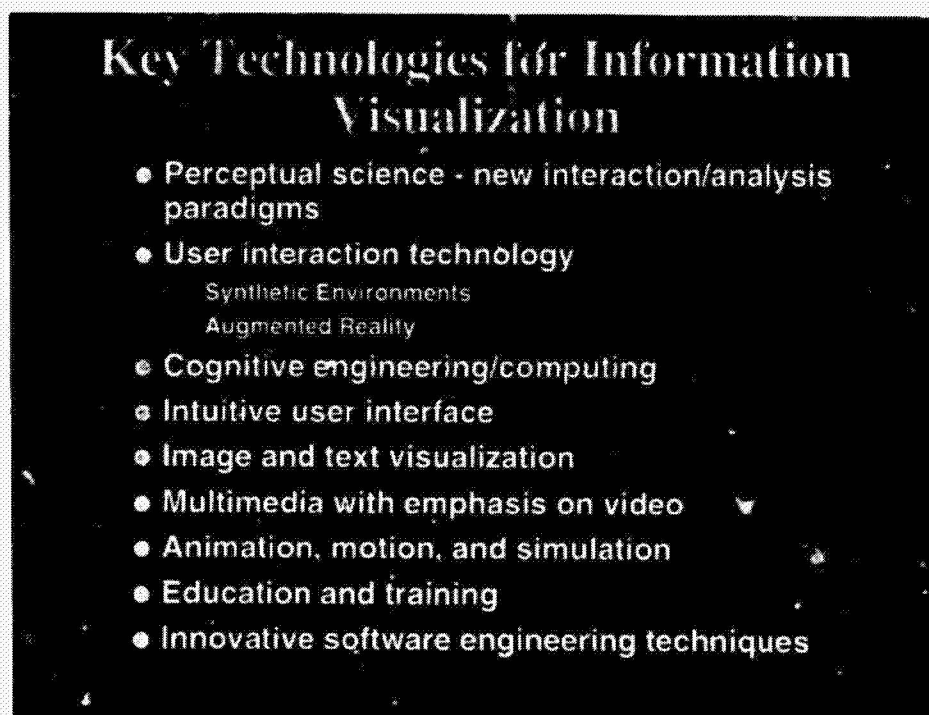
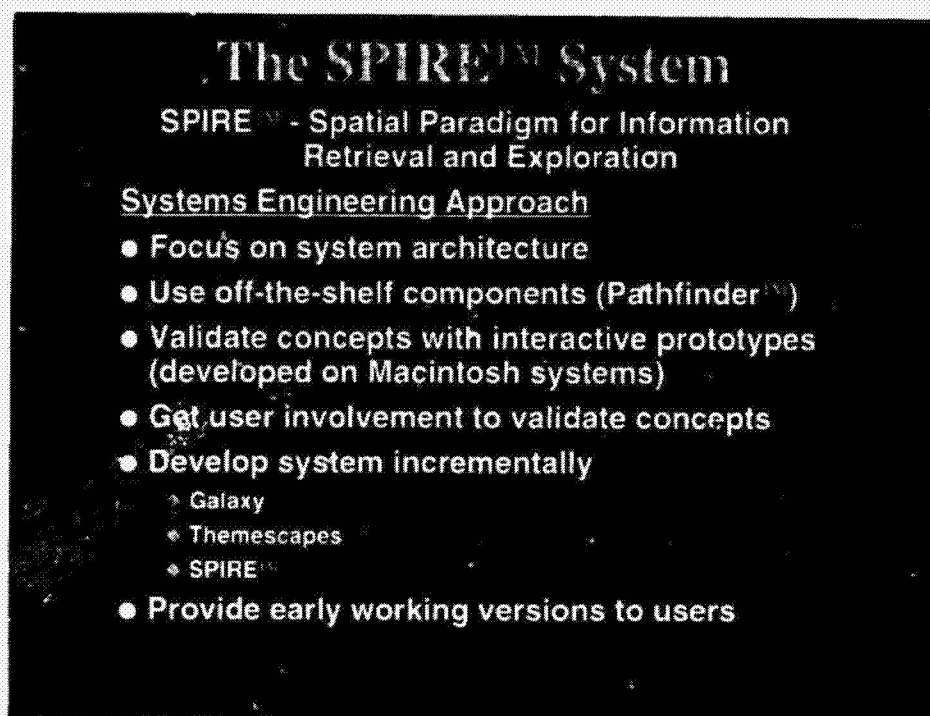


Figure 17

THE SPIRE™ SYSTEM

The work at PNL is now within a system called SPIRE - Spatial Paradigm for Information Retrieval and Exploration. The systems approach discussed below was followed.

We developed the concepts of Galaxies, received rapid feedback from users, and are now developing Themescapes within SPIRE. These will be illustrated.



The SPIRE™ System

SPIRE™ - Spatial Paradigm for Information Retrieval and Exploration

Systems Engineering Approach

- Focus on system architecture
- Use off-the-shelf components (Pathfinder™)
- Validate concepts with interactive prototypes (developed on Macintosh systems)
- Get user involvement to validate concepts
- Develop system incrementally
 - ◆ Galaxy
 - ◆ Themescapes
 - ◆ SPIRE™
- Provide early working versions to users

Figure 18

SPIRE™ SYSTEM ARCHITECTURE

This is the architecture for our technology. Note that the technology is designed to be included in other application suites. This will someday be a stand alone application, but currently must be associated with other analysis tools. Details of the architecture are contained within another paper. Please contact me if you would like more information.

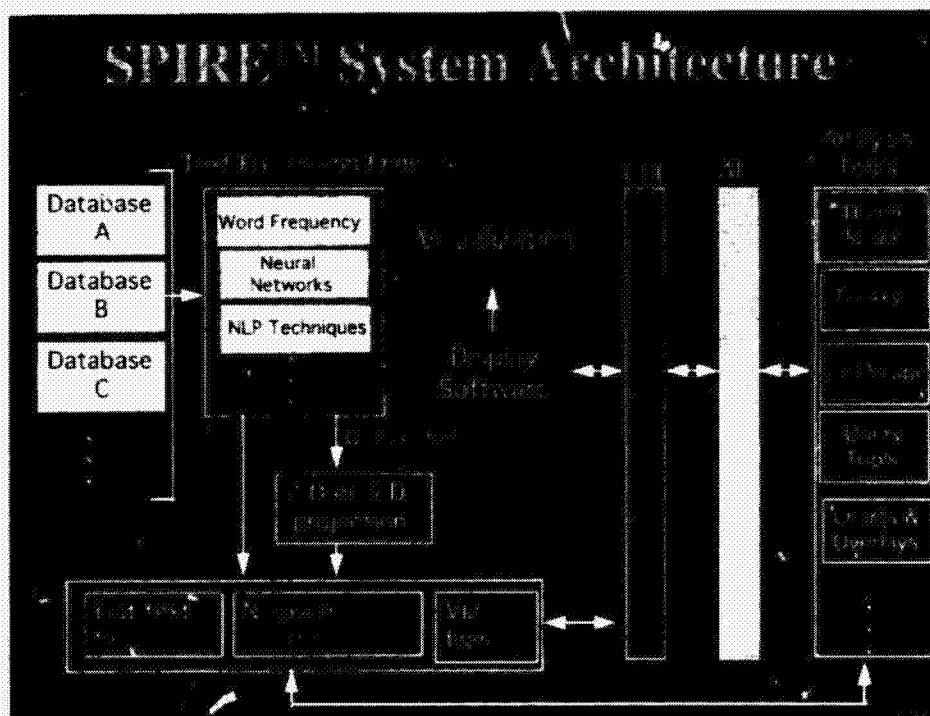


Figure 19

THE GALAXY PROTOTYPE

This is an example of the original information space from Galaxies.

The core concept is that each dot is a document. Two dots that are close to each other are similar in content. If they are far apart, then they are dissimilar in content. The clusters indicate a group of documents within similar content.

The process to gain this visual is complex. First, we obtain or calculate a proximity measure for all the words within all documents. This is a high-dimensional space. Then, we project this high-dimensional space into a view space, as illustrated.

Note that the axis means nothing. Only proximity has meaning within this visual.

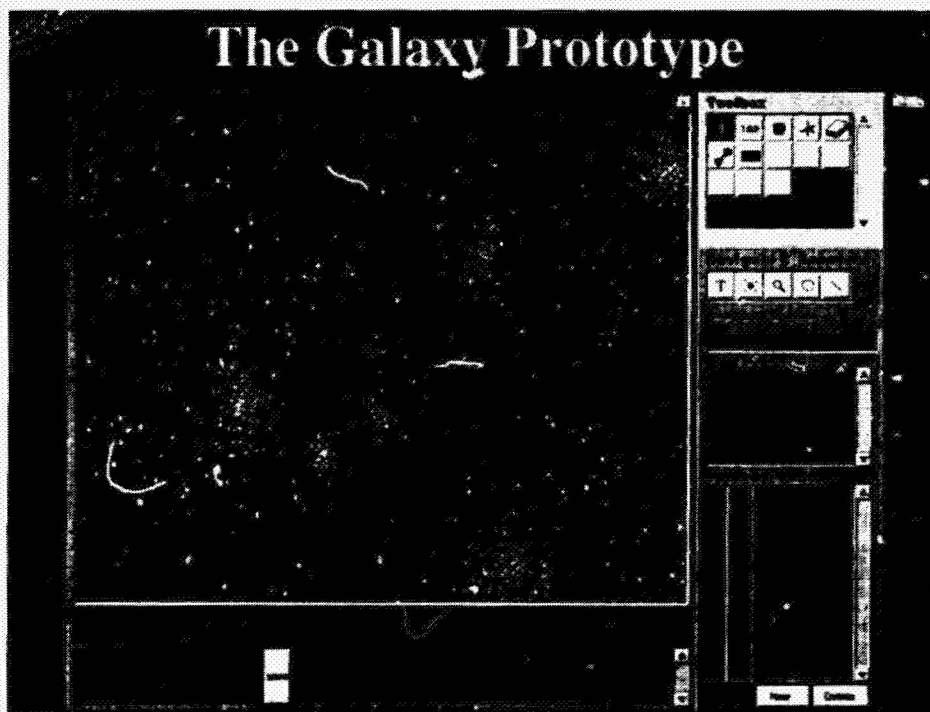


Figure 20

This is the first real analysis performed by our users. It provided the core understanding of about 600 documents so an analyst could rapidly find the right information. These were abstracts from an on-line service covering a topic. The analysis of the 600 documents was completed in less than an hour, with an understanding of what the state of this specific technology was, who was collaborating on its development, when it was developed, and what are the "close" technologies that illustrate an understanding of the issues.

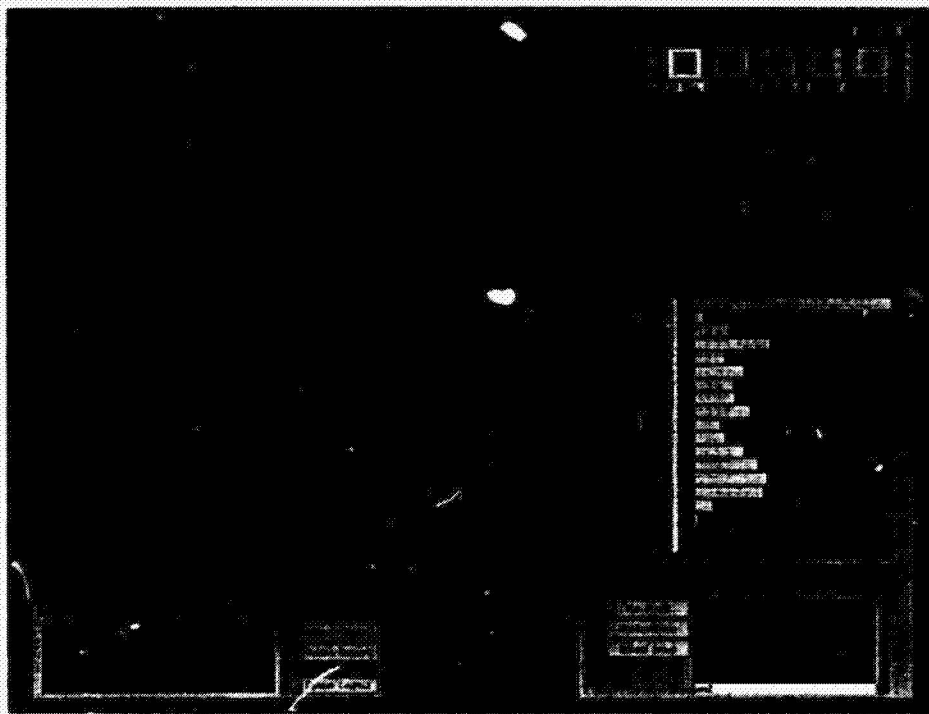
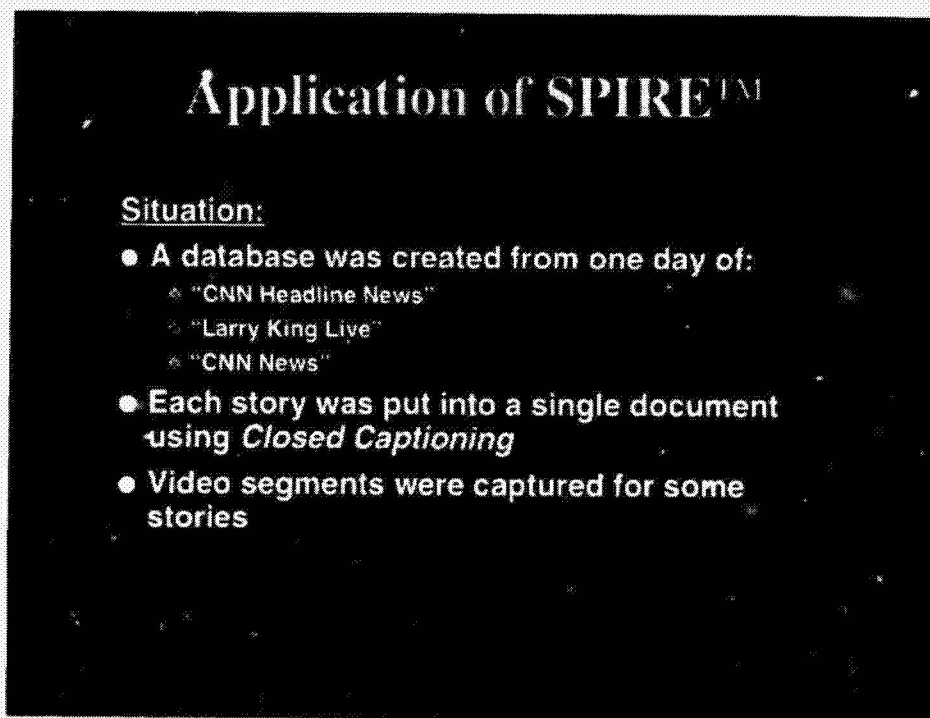


Figure 21

APPLICATION OF SPIRE™

The latest technology now allows for not only significant increase in the information space but direct visual understanding through technology now called Themescapes.

For a test data set we have selected the closed-caption from CNN. This is easily obtained through a small box connected to the television. It is important to note that this approach requires no knowledge of the topic or information space, no pre-formatting or keywording, and simply develops the proximity measures based on the content with the unstructured text. There is no human intervention in the process. This is not an AI based process.



Application of SPIRE™

Situation:

- A database was created from one day of:
 - ◆ "CNN Headline News"
 - ◆ "Larry King Live"
 - ◆ "CNN News"
- Each story was put into a single document using *Closed Captioning*
- Video segments were captured for some stories

Figure 22

This is the essence of a week's programming at CNN. The large dot represents a cluster containing N dots. Each dot is a news article. The user can open and close clusters. The proximity of one cluster to another cluster is based on the information content closeness to all others. Then within a cluster the proximity of dots is based on how each dot relates to all other clusters as well as the documents within its parent cluster.



Figure 23

We have a time slicer that allows for time-based analysis from a minute-to-a-year intervals. This illustrates a five day span of time on the CNN channel, with the last two days being highlighted.

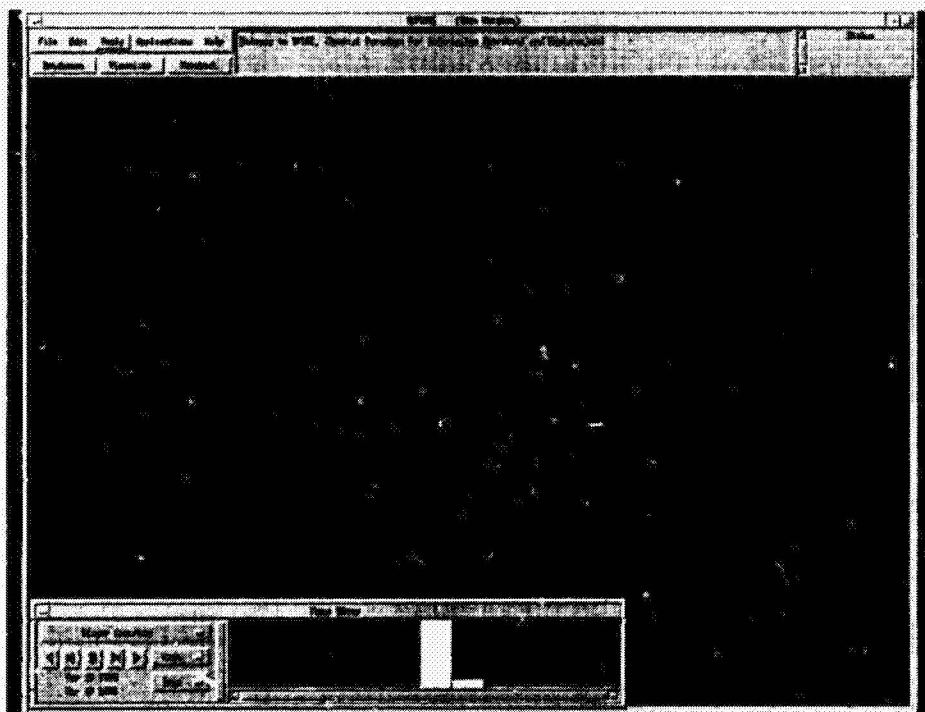


Figure 24

The lower left tool illustrates theme probing. Looking at a top view of the document clusters, one can select points, determine the primary words and make of the theme. Again, these are automatically selected based on the content and discovered relationships within the document space. The middle left tool illustrates the selection tool for subsetting, union and intersections of clusters and selected groupings via full text searches.

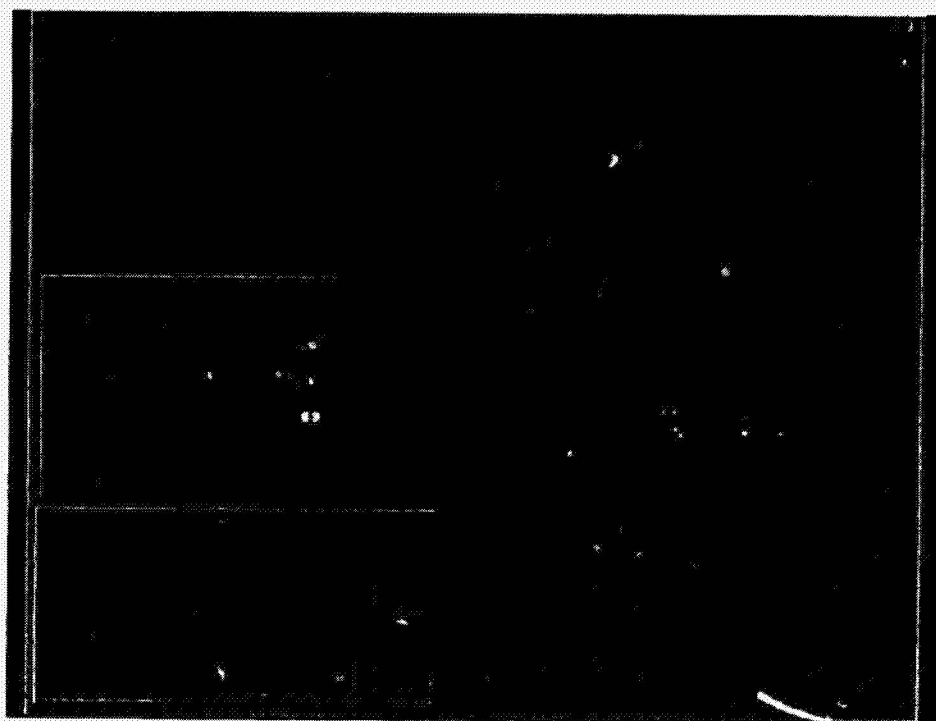


Figure 25

This visual illustrates a new concept called Themescapes. Note the basic structure similarity to the document space. This provides a landscape based on the thematic infrastructure contained within the document space. The basic principle of proximity holds. Note the themes on the right middle. These are the MCI and AT&T commercials.

The two California peaks in the lower left indicate two topics closely related. One is the weather news during the recent heavy rainstorms and the other are features about the damage.

Note the three OJ clusters. Each have a different thematic structure. They are close to the Sports clusters. The Larry King clusters are on the upper left and are separated by major topics.

Not only does the proximity contain information but the scape of the terrain is information rich.



Figure 26

This visual helps you see the shape of the themescape. The sharp slopes between the Sports and Simpson topics indicate that they are close yet there are some fundamental differences, as you might expect.

Also, note that the Shuttle is close to California as it was landing that week in California. To take full advantage of this visual, one needs a rich suite of interaction tools. A few of these tools will be available during the first release of this software.



Figure 27

One can combine and reshape all of these visuals so that the information analyst can see each one simultaneously. This is important for complex information space analysis.

This also illustrates the result of a search for all documents dealing with California.

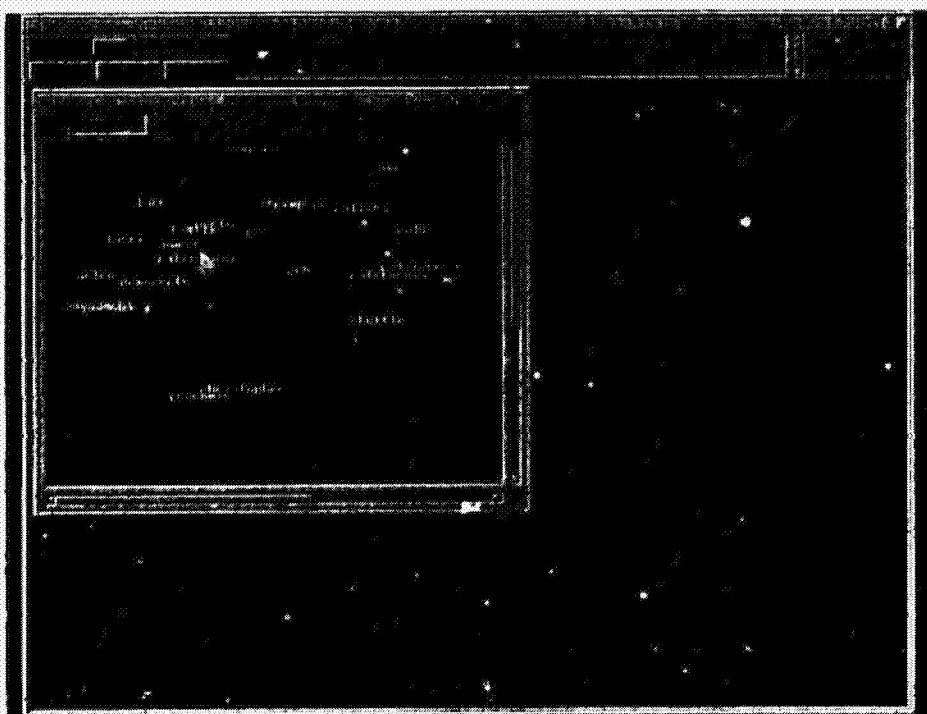


Figure 28

A video is played at this point to illustrate the dynamics of the analysis environment.

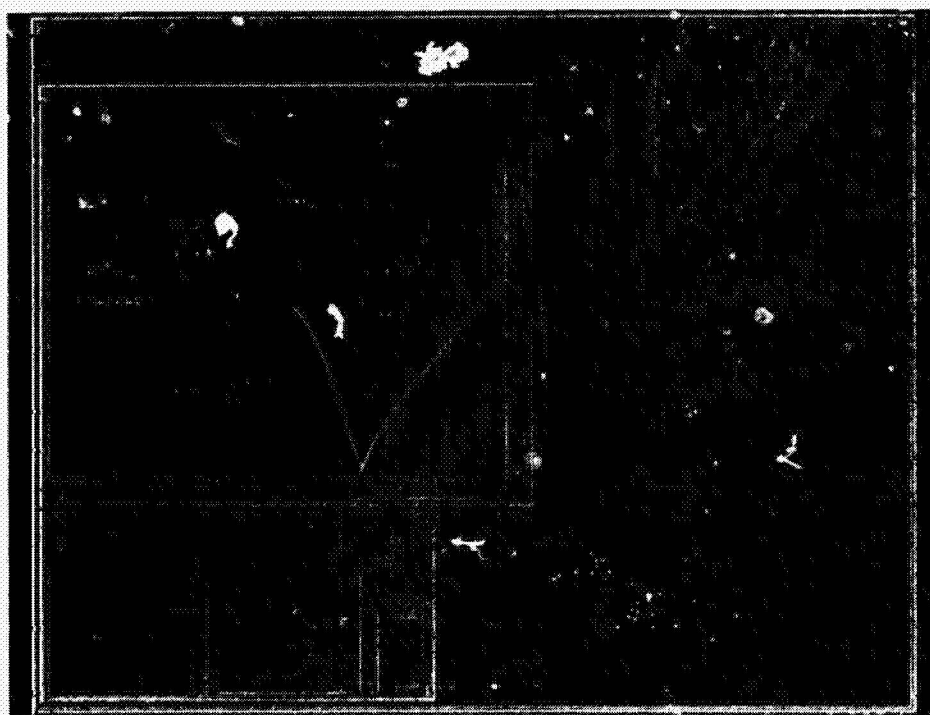


Figure 29

Also, some selected video segments were captured and can be directly played on systems that have support for video and audio.

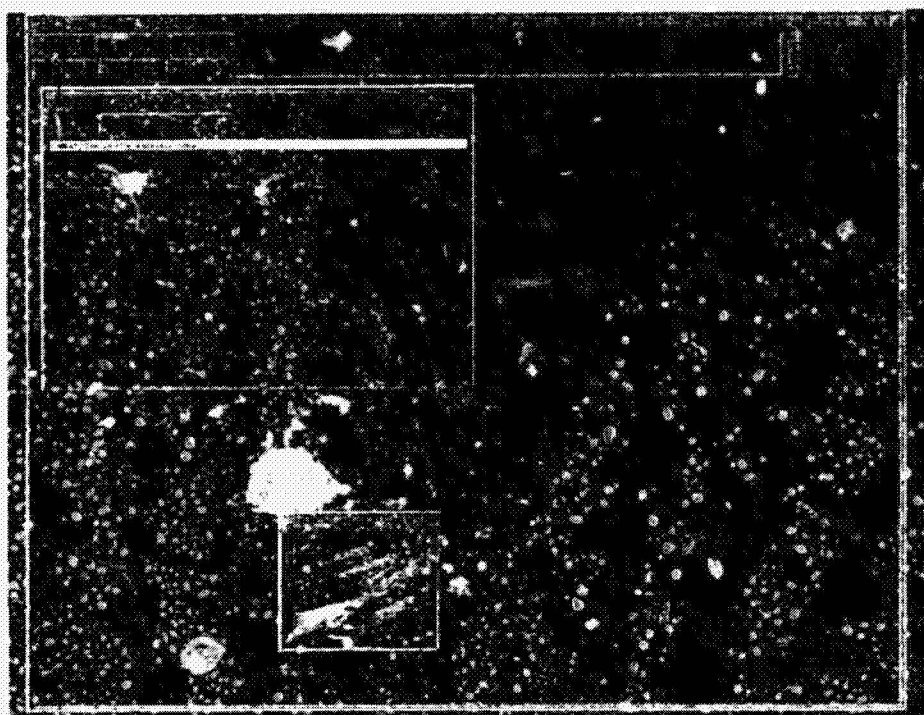


Figure 30

SPIRE™ DEVELOPMENT - KEY CONCEPTS

In summary, you have seen innovative Information Visualization that provides an analyst a new method for information analysis. This is based on a high-dimensional visual informational space that maps the information content into a spatially interactive analysis environment.

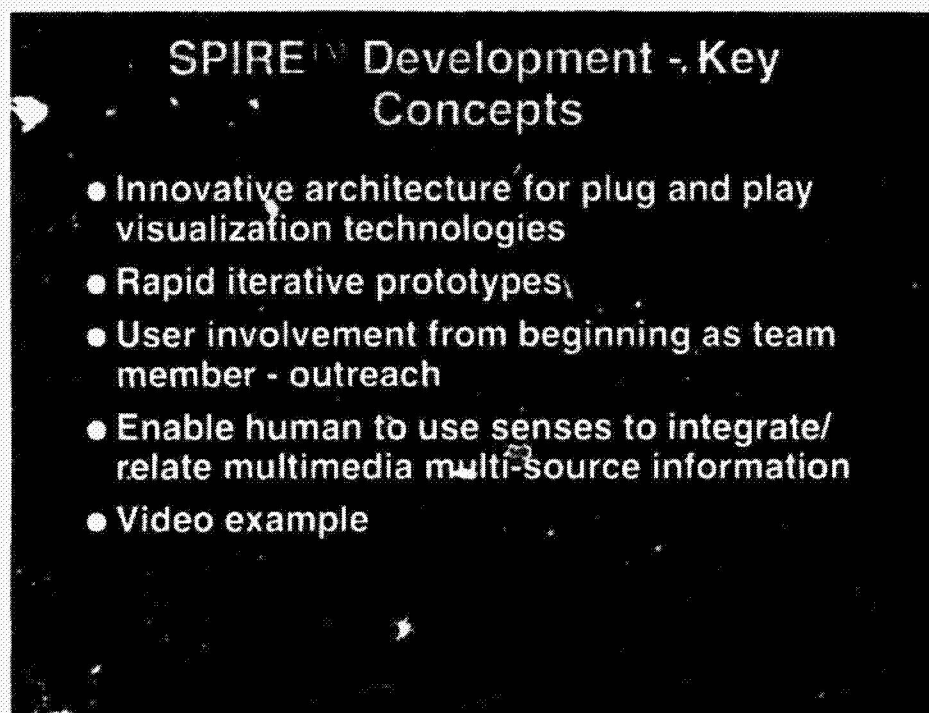


Figure 31

TAKE HOME MESSAGE FOR HCI AND COMPLEX INFORMATION SPACES

In conclusion, the key take home message is: add value by displaying all of the information by the use of high-dimensional analysis of the content spaces. The mind easily adapts to discovery of the process within these spaces.

To achieve this we need to consider all of the information spaces in selecting our fundable research tasks in this area.

**Take Home Message for HCI
and Complex Information
Spaces**

- **Include all information resources**
 - text, images, diagrams, procedures, regulations
 - market information
 - quality information
 - materials information
- **Add value through information visualization
and the HCI**

**Focus R&D on total information spaces
for today's and tomorrow's industry**

Figure 32

**NEXT
DOCUMENT**

5.7-54
C252

THE POWER OF PEZONOMICS

129893

Joel Orr
Orr Associates International
Chesapeake, Virginia

There are not many virtual worlds in working environments today, but the few that there are suffer from a flaw similar to the biggest fault of less-engaging computer systems: they are designed with little consideration for the way people really are, and the way they behave. They are, in a word, ergonomically unsound.

Inconsistency is one of the biggest problems. We have grown to expect different programs to have different ways of doing things. Type ALT-F, X to exit from one program, CONTROL-Q to leave another. If cars were as inconsistent, renting one in a strange city could be a problem: "This new ergonomic model has the brake pedal on the right, and the accelerator is operated with your left elbow. Just initial in these eleven places, and sign here—sir? Where'd he go?" Out to the taxi stand.

Disregard for the human shape is another. Take the standard computer keyboard. The arrangement of its keys was designed in the nineteenth century to slow down typists, to keep them from typing another letter before the previous letter key had a chance to fall away from the platen of the typewriter. Moreover, the keyboard's size has been reduced, while hand sizes have grown. And the keys are not positioned to meet naturally-placed human fingers. That successful design is still working, sending the children of hundreds of orthopedists through college by creating a whole branch of the repetitive stress injury market.

Indifference to health and safety is not unusual among designers of hardware and software. Manufacturers reduce harmful radiation from CRT's only under duress, and they do not want to know what electromagnetic fields might be doing to people.

To add insult to injury, most computer systems are just not *fun* to learn or to use; they are arcane drudgery.

We need to think about human interfaces to computer systems in general, and virtual worlds in particular, in a new way; a good place to start the revolution is with terminology. *Pezonomics* is my replacement for ergonomics. Ergonomics comes from two Greek words that mean "work" and "law"—"the law of work," in other words. Too austere! The most popular, easy-to-learn computer systems in the world are video games, prime examples of successful application of pezonomics—"the law of play." We need to capture the essence of play and calibrate our computer systems to its cadences.

Why do we love to play? Perhaps it is that in play, we have permission to win or lose; we find closure; and we can exercise skills in a protected environment.

The places in which we try out new things—new designs, new concepts—ought to somehow be environments in which consequences are attenuated. That is the power of simulation: the consequences of crashing a Boeing 747 simulator are noticeably milder than those of crashing a real 747. So we can make mistakes, and live to learn from them.

This is one of the great promises of VR. A virtual world has *controlled* consequences, but at the same time, it engages and involves the participant in active ways; she or he is not merely an observer.

When introducing people to new things, I measure their responses on a scale of *astonishment*. The unit of this scale is the *gasp*, when positive, and the *yawn*, when negative. Extremes are not desirable; a five-gasp novelty could lead to heart failure, while a five-yawn novelty may result in snoring.

A friend once told me he designs systems according to the “principle of minimal astonishment.” He said his goal was to minimize surprises to his customers, both in the creation of the system and in its use. “The customer or user should always know what is going to happen next, and should have some control over that next event,” said this wise man.

But as in marriage, complete predictability can lead to boredom. To maintain interest, we need *controlled surprise*; flowers and romantic greeting cards without a birthday or anniversary, unplanned acts of warmth and kindness. Good programs and virtual worlds should have an unending stream of pleasant surprises, nice things that happen to the user when they are not expected.

Perhaps the simplest approach to pezonomics involves engaging the “right brain” through the use of graphics. Normal people find computer systems unapproachable. Graphical user interfaces help—by providing spatial access to sequential pieces, much like what writing does for speechwriting, and by offering the user a consistent approach to multiple applications, thereby reducing the cost of trying out a new application. Mac users, Apple reported a few years ago, use six to eight applications on average, while DOS users employed only two. A more recent study shows that Windows users have caught up with Mac users. GUIs have pezonomic qualities.

Rhythm is important in pezonomics. Human beings are complex systems—too complex to have simple natural frequencies. But there are certain frequencies that “resonate”—or at least, interact—with some human phenomena. Low-frequency sound pulses at or near a person’s heart rate seem to cause the human system to “lock in” to the frequency of the sound generator; once this occurs, changes in the frequency cause corresponding changes in the person’s heart rate, as well as in other physical functions. The most popular video games are not the ones with the best graphics; they are the ones that have an audible heartbeat-rate low-frequency pulse, that accelerates as the game progresses. This auditory entrainment causes the player’s heart rate to speed up, with an accompanying production of adrenaline and endorphins. By the end of the game, the player is “hyped”—and wants more. (My term for the rhythmic aspects of the person-computer interface is *anthropocybersynchronicity*.)

This entrainment also creates a deep sense of *rapport*—special connection. Good sales people have long known what practitioners of neurolinguistic programming have recently written about: you can establish rapport with someone by intentionally mirroring different aspects of their behavior—their rate of breathing, their blinking rate, the rate at which their leg is swinging, for example. After a few minutes of matching, you can verify that you have rapport by leading—changing the rhythm, and watching to see if they follow. If they do, you are communicating with the person on a very primal level, and they are much more open to your suggestions and other forms of leading than when such rapport is absent.

Studies of people in singles bars back this up. Observers noted that people who began to mirror each other’s behavior soon left together; people who were “out of synch” with each other after a few minutes separated and made other contacts.

Proper pezonomics will also result in increased person-computer *coupling*. The computer is a general-purpose tool, something we use to do a job. We must measure its effectiveness by how easily and how well it helps us accomplish our goal—which is usually *not* operating the computer, it is writing, accounting, designing, drafting, or something to which the computer—except for the specifics of its assistance—is irrelevant.

We can increase our control of the tool by increasing our coupling to it—the extent to which our actions and the actions of the computer system affect each other. Coupling is a kind of rapport; rhythm, through resonance, enables us to increase that coupling.

Of course, increasing coupling might give the tool more control over the user, which could be undesirable; like the binding of a ski, it has to be both loose and tight. You do not want the ski to fall off while you are going down a slope, but you want it to come off easily if you fall.

Our goal is to get at our work, our art, our play; the computer is only a means to an end. It is like reading, or learning to understand pictorial imagery. Ultimately, it is not literacy, or pictoracy, we desire; it is not even “mediacy,” a facility with multimedia. Rather, we yearn for *immediacy*—enhanced access to our problems so that we are empowered to solve them without apparent mediation, without the intrusion of the irrelevancies of the computer or our disabilities. Pezonomics can bring us closer to this goal.

Suppose the headband of your head-mounted display had a sensor that could detect your pulse. When you first put it on, the system would ask you about your mood and alertness, from time to time. It would build a table with the corresponding heart rates. After a period of calibration, it could then sense your level of alertness, and use rhythmic auditory and visual pulsation to alter it. It could, for example, flash a feature in your field of view, subtly, at the rate of your pulse, while making an unobtrusive but audible clicking sound. When it detected synchronization between your pulse rate and its beat, it could speed up the flashing and clicking, checking to see that your pulse was entrained.

By increasing your heart rate, the system would cause your body to generate the substances that are the concomitants of the “fight or flight” response—including endorphins and enkephalins, pain-blunting, pleasure-enhancing, morphine-like chemicals that could make you more effective.

They could also make you less effective, if your work required a more contemplative mood. And there is now evidence that activation of the “fight or flight” response causes stress and disease. For this reason, you had better be able to control what the system does to you.

“What do you call people who practice the rhythm method of birth control?” goes the riddle. “Parents,” is the answer. Like any other tool or approach, rhythm does not ensure success.

Rhythm can be a powerful ally or a formidable foe, a liberator or an enslaver. I do not believe it is intrinsically evil, but it can be used for evil purposes, such as controlling people against their will. We should approach it cautiously and intelligently, respecting its destructive power while we harness it for our benefit.

We trustingly submit our entire sensoria to immersive virtual reality. Yet participants in virtual worlds can become disoriented, even to the point of nausea or more enduring ill effects. Perhaps these can be mitigated through pezonomic rhythm management.

System designers seek better person/machine coupling. To date, we have looked at human interface issues with the mathematics of Euclid and Newton, whose underlying assumptions derive from Plato's: everything in the world is an approximation of an ideal. Recent discoveries in “the mathematics of chaos” reveal that things are both simpler and more complex than we ever imagined. Sealed mysteries of natural phenomena, and biological phenomena in particular, are yielding in embarrassing profusion to this new “Open Sesame.” I wonder if there will be found, in a biological setting, a chaotic or fractal analog to the Newtonian notion of resonance, one that will help us design virtual worlds that will liberate people without endangering them.

Conclusion: Pezonomics—the law of play—must be developed and applied to computer systems in general, and VR in particular, if these artificial environments are to become hospitable for humans. Without playful interfaces, virtual worlds will be as arcane and inaccessible as UNIX. The coming information superhighway promises a high-bandwidth information infrastructure that will make possible a shared multisensory space. We must pad its walls pezonometrically or endanger the innocent.

The Power Pezonomics

*Safe and Engaging
Immersion through
Rhythm and Play*

■ Joel N. Orr
President
Virtual Worlds Society

oooooooo

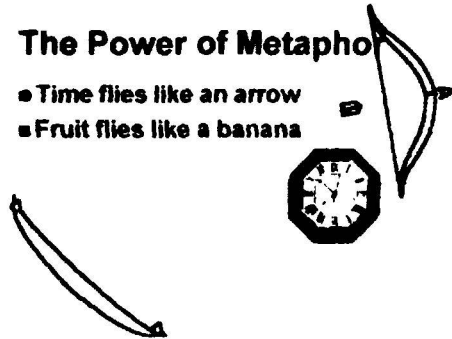
VR: A Dangerous Place?

- Sensory "amputation"
- Disorientation
- False signals
- Electrical & mechanical effects
- Physiological problems
- Psychic challenges

oooooooo

The Power of Metapho

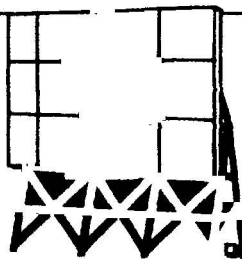
- Time flies like an arrow
- Fruit flies like a banana



oooooooo

*Between the idea
And the reality
Between the notion
And the act
Falls the shadow*

--from T. S. Eliot, "The Hollow Men"



oooooooo

The Fertile Verge

- Increased sense of community
- Greater self-awareness
- More tolerance for innovation

oooooooo

What is it?

- Artificial reality
- God-like power in a forgiving world
- "Transportation system for the senses"
- Immersion
- Immediacy
- Intimacy



oooooooo

"Every man, wherever he goes,
is surrounded by
a cloud of comforting convictions,
which follow him
like flies on a summer day."

Bertrand Russell

oooooooo

Any medium has the power
to restructure our minds in a
unique way by imposing its
own mode of thought.

Marshall McLuhan

oooooooo

Functions of VR

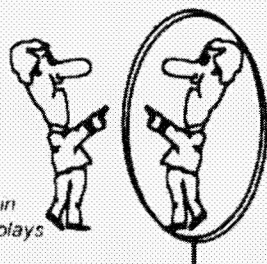
- Immersion
- Simulation
- Buffering



oooooooo

Functions of VR

- Mirroring



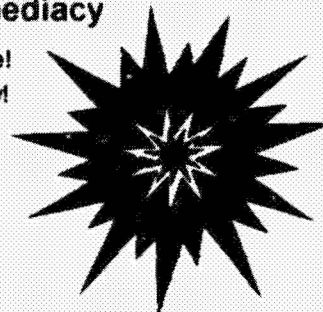
Behavior is a mirror in
which every one displays
his image

Goethe

oooooooo

Immediacy

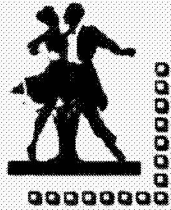
- Here!
- Now!



oooooooo

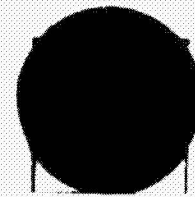
Intimacy

- Negroponte's prophecy
- Richard Avedon's "unearned..."
- The cognitive definition



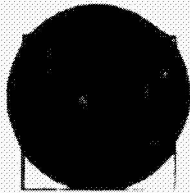
The Human/Computer Interface

Ergonomics?



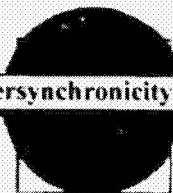
The Human/Computer Interface

Pezonomics?



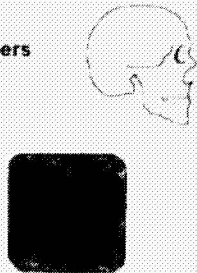
The Human/Computer Interface

Anthropocybersynchronicity!



The Dark Side

- Electromagnetic dangers
- Ergonomic issues
- Loss of reality
- Loss of control
- Loss of humanity



Effects

People using telephone, TV, or radio, or any other electric service do not have bodies; they are discarnate beings.



Marshall McLuhan

More Warnings from the Prophet

"We recognized that communication satellites had ushered in a new Global Electric Theatre of the Absurd, where only the unexpected happens, precisely because old logical Thinking can no longer keep pace with new eco-logical Being."

"Electric speedup had created global crises of identity, regardless of any intent to improve human communication."

oooooooo

What Should We Do? ? ?

- Add rhythm to the research agenda
- Focus on play
- Avoid deconstructionism



oooooooo

In Short...

- VR may well be "the promised land," but
- with its milk and honey, there are "giants" in the land:
- overcoming them is possible, but
- requires planning.
- Think "whole"!



Questions?

Write or call me:

75300 1514@compuserve.com

804/467-2677; fax 804/495-8548

oooooooo

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE November 1995	3. REPORT TYPE AND DATES COVERED Conference Publication		
4. TITLE AND SUBTITLE Human-Computer Interaction and Virtual Environments		5. FUNDING NUMBERS WU 505-63-50-17		
6. AUTHOR(S) Ahmed K. Noor, Compiler				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) NASA Langley Research Center Hampton, VA 23681-0001 and University of Virginia Center for Computational Structures Technology Hampton, VA 23681-0001		8. PERFORMING ORGANIZATION REPORT NUMBER L-17547		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Washington, DC 20546-0001		10. SPONSORING/MONITORING AGENCY REPORT NUMBER NASA CP-3320		
11. SUPPLEMENTARY NOTES Ahmed K. Noor University of Virginia Center for Computational Structures Technology Hampton, VA 23681-0001				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified-Unlimited Subject Category 62		12b. DISTRIBUTION CODE		
13. ABSTRACT (Maximum 200 words) This document contains presentations and discussions from the joint UVA/NASA Workshop on Human-Computer Interaction and Virtual Environments held at the location of the Virginia Consortium of Engineering and Science Universities in Hampton, VA on April 26-27, 1995. The presentations addressed activities in the areas of human factors, virtual environment, sonification, haptic interfaces, distributed and collaborative environments, and engineering applications of some prototype synthetic environments. Workshop attendees represented NASA, NIST, WPAFB, Sandia National Labs, automotive and commercial software industries, and academia. The workshop objectives were to assess the state-of-technology and level of maturity of several human-computer interaction facilities, their potential applications to future engineering systems; and to provide directions for future research.				
14. SUBJECT TERMS Synthetic environment; Virtual environment; Human factors; Telerobotic systems; Sonification; Auditory display; Haptic interfaces; Distributed and collaborative environments			15. NUMBER OF PAGES 315	
			16. PRICE CODE A14	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	